



US008781727B1

(12) **United States Patent**  
**Bonawitz et al.**

(10) **Patent No.:** **US 8,781,727 B1**  
(45) **Date of Patent:** **Jul. 15, 2014**

(54) **METHODS AND SYSTEMS FOR PERFORMING FLOCKING WHILE EXECUTING A LONG-RANGE FLEET PLAN**

(71) Applicant: **Google Inc.**, Mountain View, CA (US)

(72) Inventors: **Keith Allen Bonawitz**, Berkeley, CA (US); **Dan Piponi**, Oakland, CA (US)

(73) Assignee: **Google Inc.**, Mountain View, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **13/742,209**

(22) Filed: **Jan. 15, 2013**

(51) **Int. Cl.**  
**G01C 21/12** (2006.01)

(52) **U.S. Cl.**  
USPC ..... **701/422; 701/410**

(58) **Field of Classification Search**  
USPC ..... **701/422, 410**  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

4,802,096	A *	1/1989	Hainsworth et al. ....	701/301
6,167,263	A	12/2000	Campbell	
6,324,398	B1	11/2001	Lanzerotti et al.	
6,402,090	B1	6/2002	Aaron	
6,437,743	B1	8/2002	Mintz et al.	
6,507,739	B1	1/2003	Gross et al.	
6,611,755	B1 *	8/2003	Coffee et al. ....	701/482
6,628,941	B2	9/2003	Knoblach et al.	
6,678,520	B1	1/2004	Wang	
6,694,235	B2 *	2/2004	Akiyama .....	701/29.6
6,799,100	B2 *	9/2004	Burns et al. ....	701/25
6,842,674	B2 *	1/2005	Solomon .....	701/23
7,046,934	B2	5/2006	Badesha et al.	
7,356,390	B2	4/2008	Knoblach et al.	

7,469,857	B2	12/2008	Voss	
7,512,462	B2	3/2009	Nichols et al.	
7,844,218	B2	11/2010	Seligsohn et al.	
7,946,533	B2	5/2011	Goodzeit	
8,116,763	B1	2/2012	Olsen	
8,170,747	B2 *	5/2012	Chen et al. ....	701/36
8,234,067	B2 *	7/2012	Bauer et al. ....	701/467
8,238,903	B2	8/2012	Korb et al.	
8,260,485	B1 *	9/2012	Meuth et al. ....	701/26
8,275,499	B2	9/2012	Coulmeau et al.	
8,494,689	B1 *	7/2013	Ashton .....	701/3

(Continued)

**FOREIGN PATENT DOCUMENTS**

FR 2972697 9/2012

**OTHER PUBLICATIONS**

Mike Hanlon, "New Software allows a flock of UAV's to work together". Mar. 28, 2005; Gizmag.com.\*

(Continued)

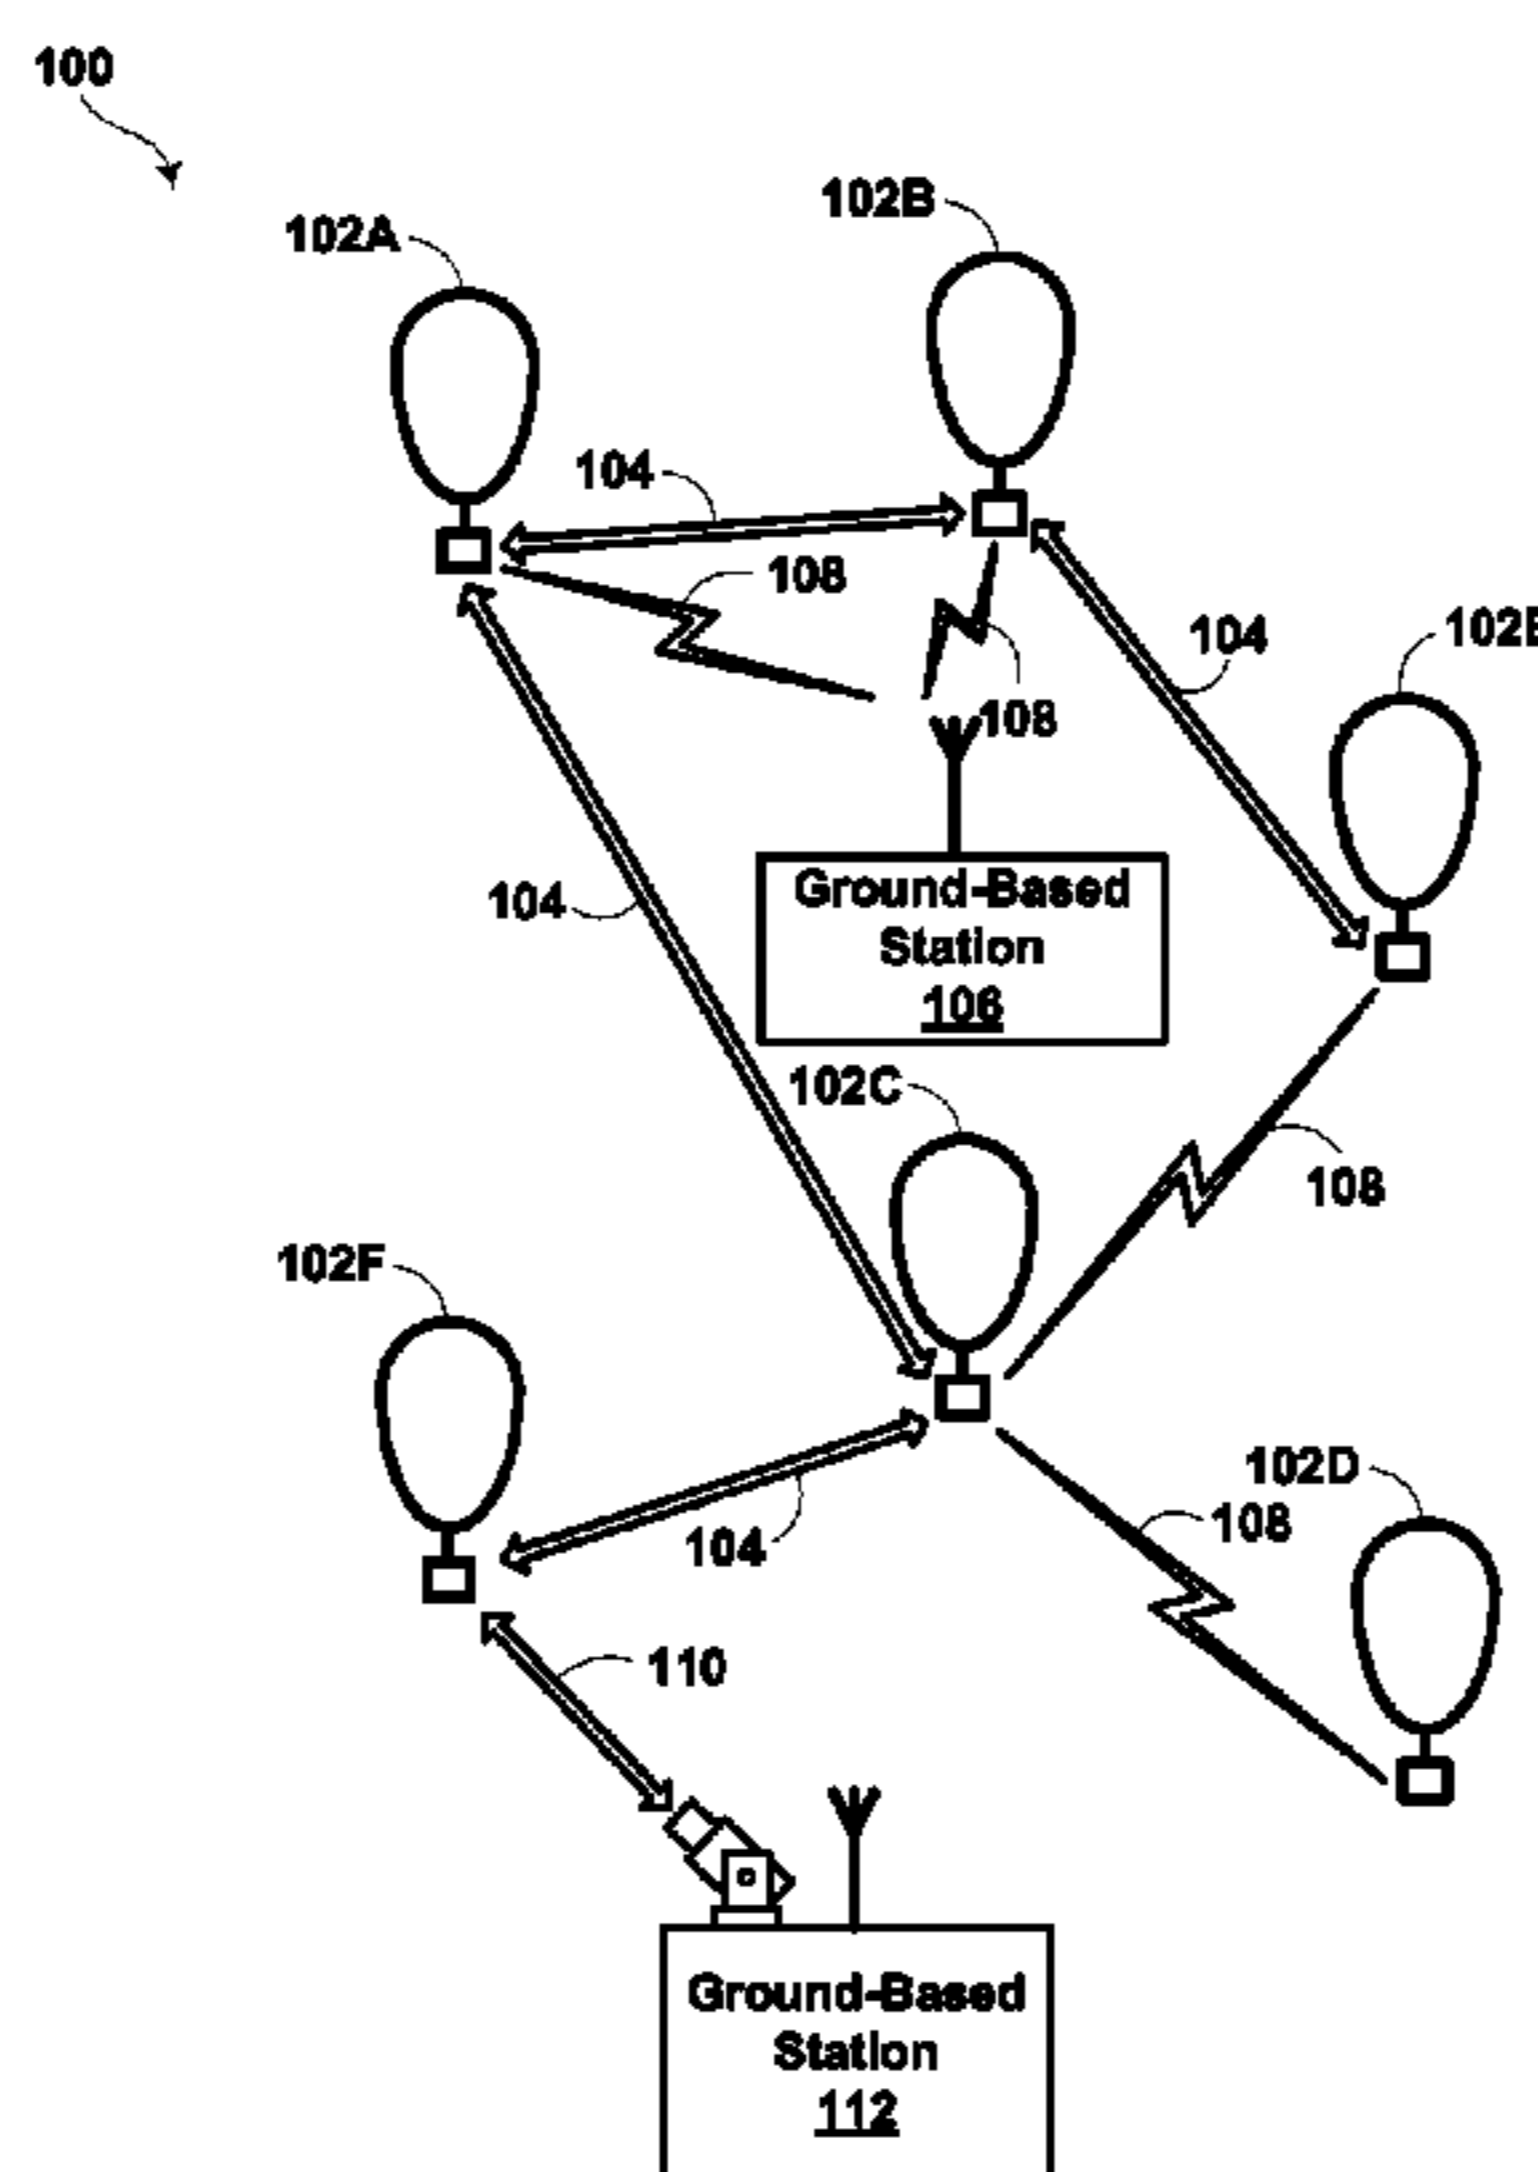
*Primary Examiner* — Thomas Black  
*Assistant Examiner* — Tyler Paige

(74) *Attorney, Agent, or Firm* — McDonnell Boehnen Hulbert & Berghoff

(57) **ABSTRACT**

Methods and systems for performing flocking while executing a fleet plan are provided. An example method includes receiving a sequence of coverage requirements for a region and an associated period of time, and determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time. Additionally, based on a determined sequence of intended destinations for a vehicle of the one or more vehicles and based on a desired spatial relationship between the vehicle and one or more neighboring vehicles for a given time period, a flocking-based direction of travel for the vehicle may be determined for the given time period.

**20 Claims, 14 Drawing Sheets**



(56)

**References Cited**

U.S. PATENT DOCUMENTS

2002/0077944	A1	6/2002	Bly et al.	
2003/0040273	A1	2/2003	Seligsohn et al.	
2003/0141409	A1*	7/2003	Lisoski et al.	244/13
2006/0167599	A1	7/2006	Bodin et al.	
2007/0288132	A1*	12/2007	Lam	701/23
2008/0144884	A1	6/2008	Habibi	
2009/0099768	A1*	4/2009	Bauer et al.	701/206
2009/0267740	A1	10/2009	Pizzuto	
2011/0147513	A1	6/2011	Surmont	
2012/0073682	A1	3/2012	Geneste	
2012/0215505	A1	8/2012	Srivastav	
2012/0256770	A1	10/2012	Mitchell	

OTHER PUBLICATIONS

“Fleet management with automatic vehicle location”, by Harden, M.D., published May 20-22, 1986 (Motorola Inc., Mobile Products Division).

“A Mobile Location—Based Vehicle Fleet Management Service Application”, by Silva, A.P., published Jun. 9-11, 2003, IEEE.

A method for Balloon Trajectory Control, by Aaron, K.M. et al., Global Aerospace Corporation, published Jun. 21, 2008.

Path Planning for Autonomous Underwater Vehicles in Realistic Oceanic Current Fields: Application to Gliders in the Western Mediterranean Sea, by Garau B. et al., Journal of Maritime Research, vol. VI. No. II, pp. 5-22, 2009.

Real-time Trajectory Design for Unmanned Aerial Vehicles using Receding Horizon Control, by Yoshiaki Kuwata, Thesis submitted to the MIT Department of Aeronautics and Astronautics, Jun. 2003.

Swarm Intelligence in autonomous collective robotics: from tools to the analysis and synthesis of distributed control strategies, Alcherio Martinoli, Thesis, 1999.

Autonomous Underwater Vehicle Path Planning: A Comparison of A and Fast Marching Method, Kevin DeMarco, available at [http://www.kevindemarco.com/wp-content/uploads/2011/05/fmm\\_paper\\_revised.pdf](http://www.kevindemarco.com/wp-content/uploads/2011/05/fmm_paper_revised.pdf), 2011.

Optimal AUV path planning for extended missions in complex, fast-flowing estuarine environments, Dov Kruger et al., 2007 IEEE Conference on Robotics and Automation, Apr. 2007.

\* cited by examiner

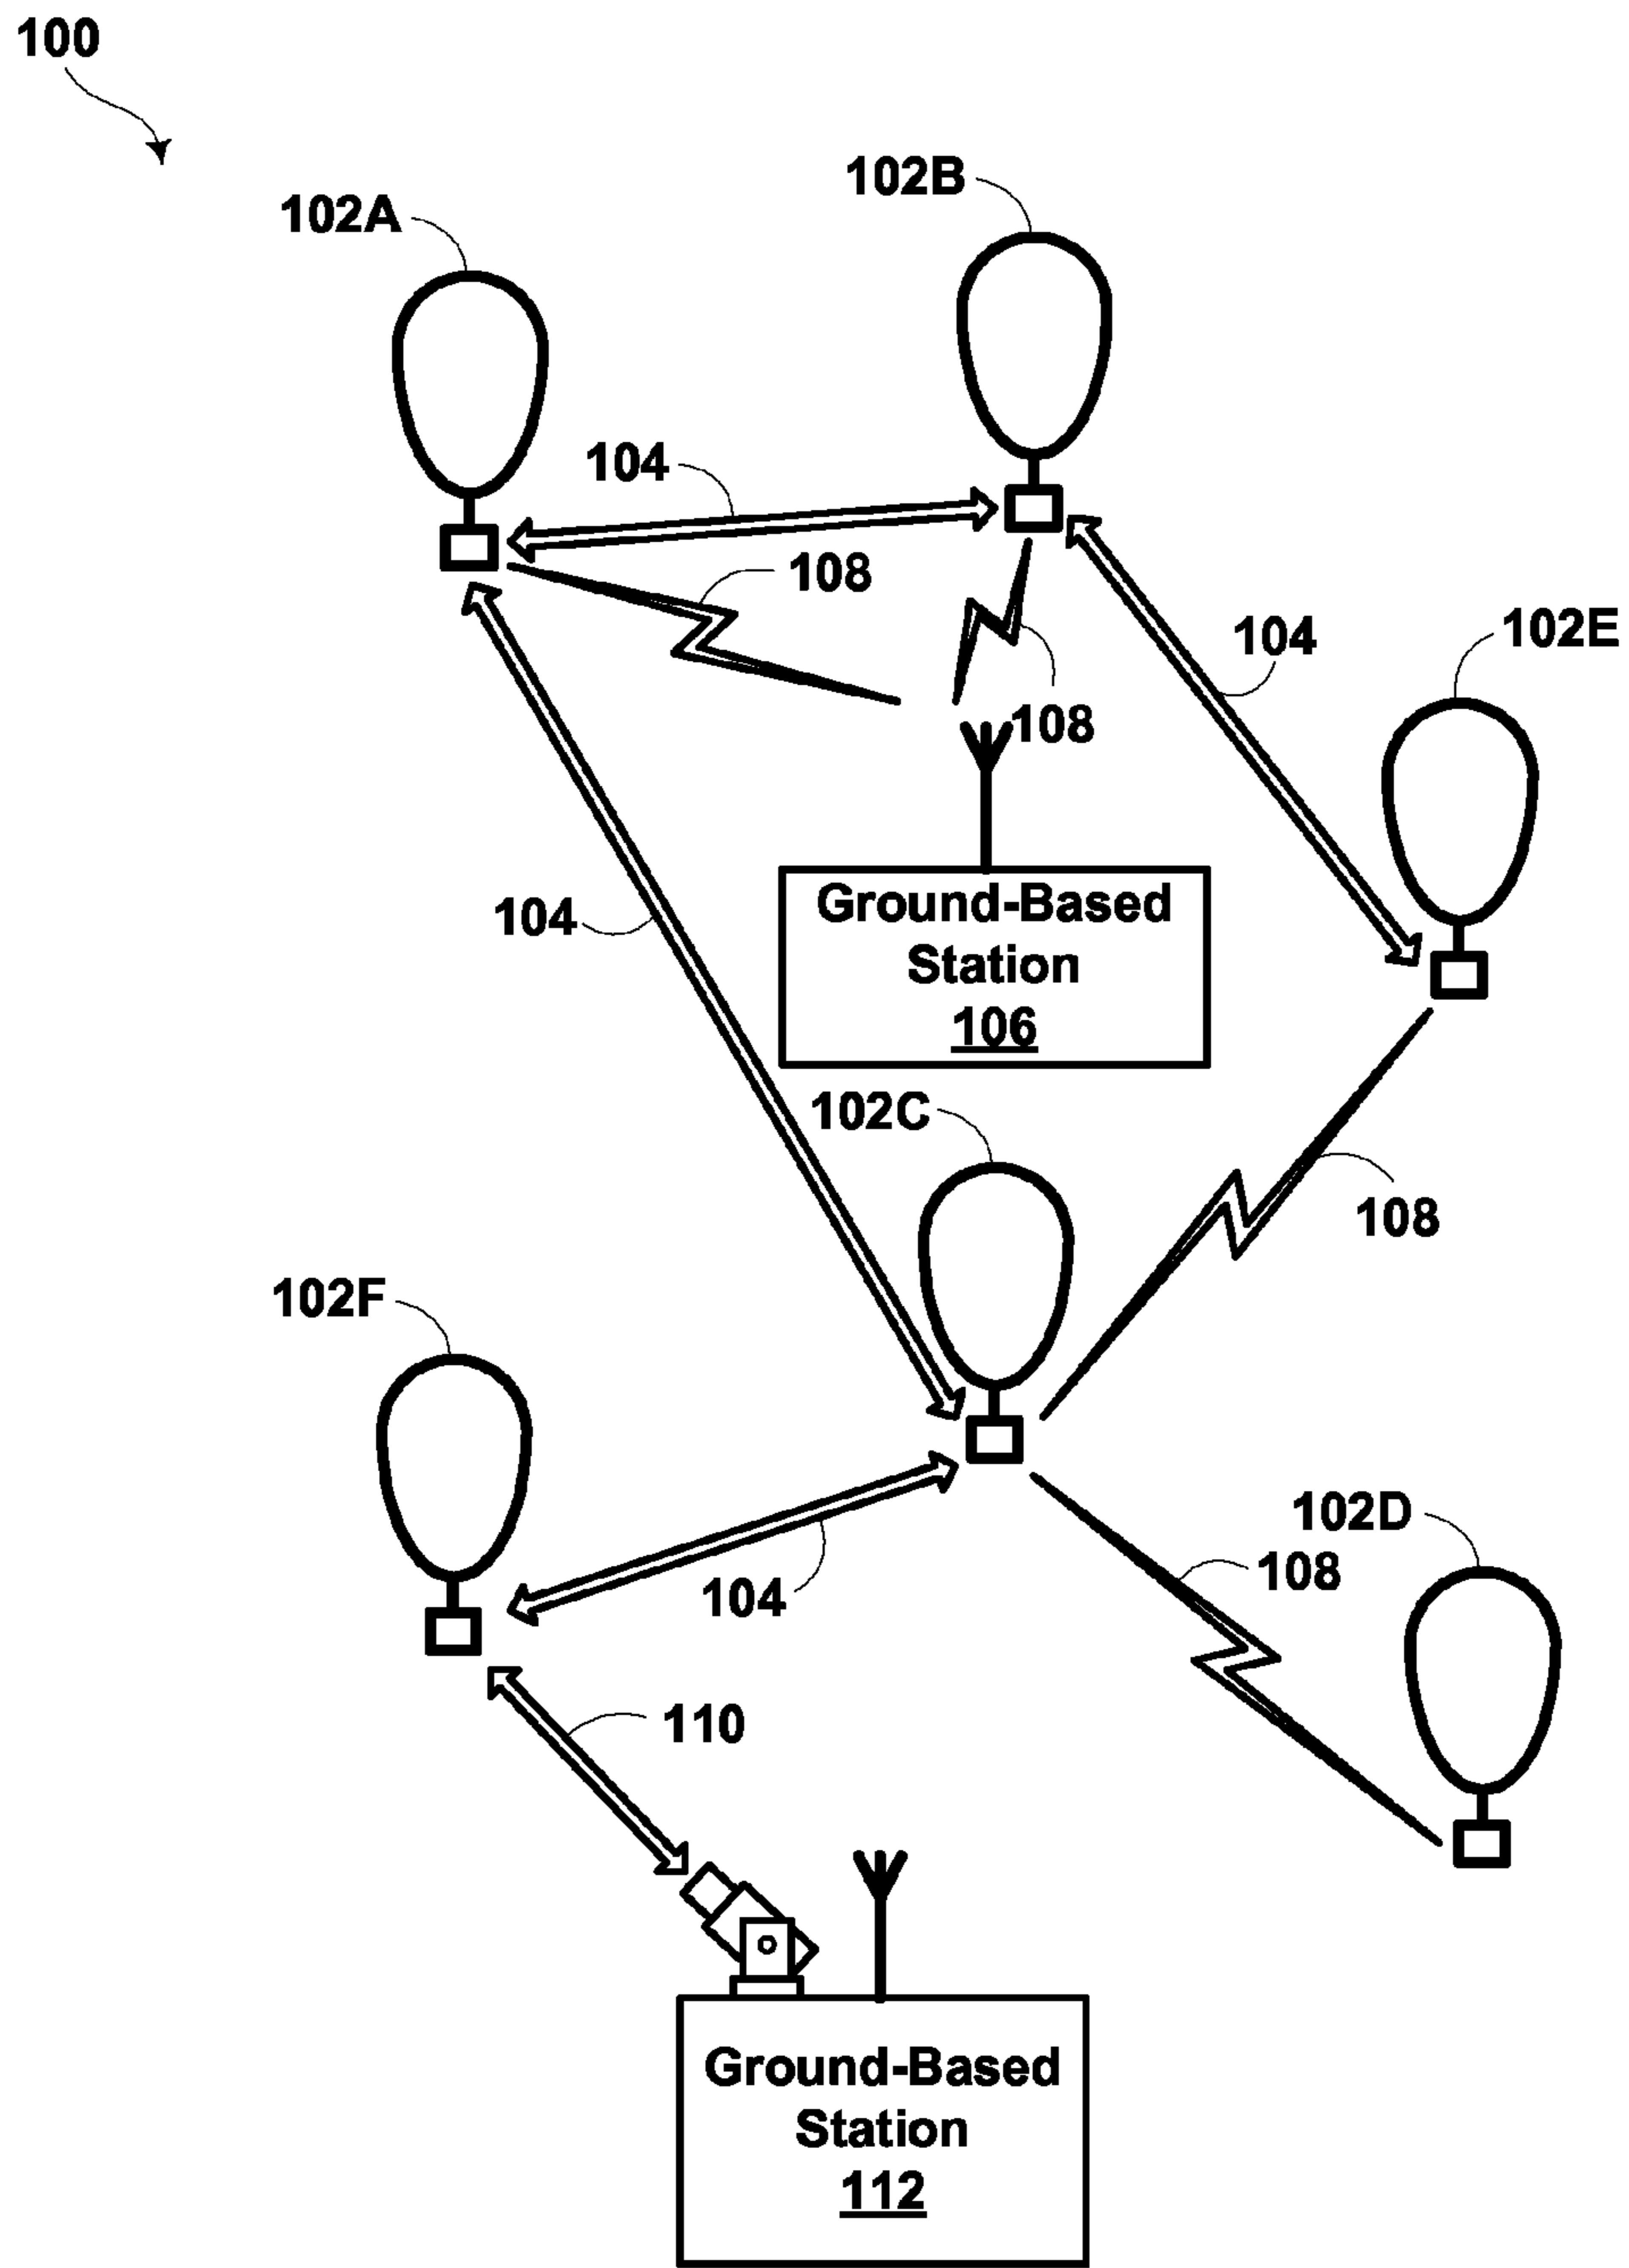
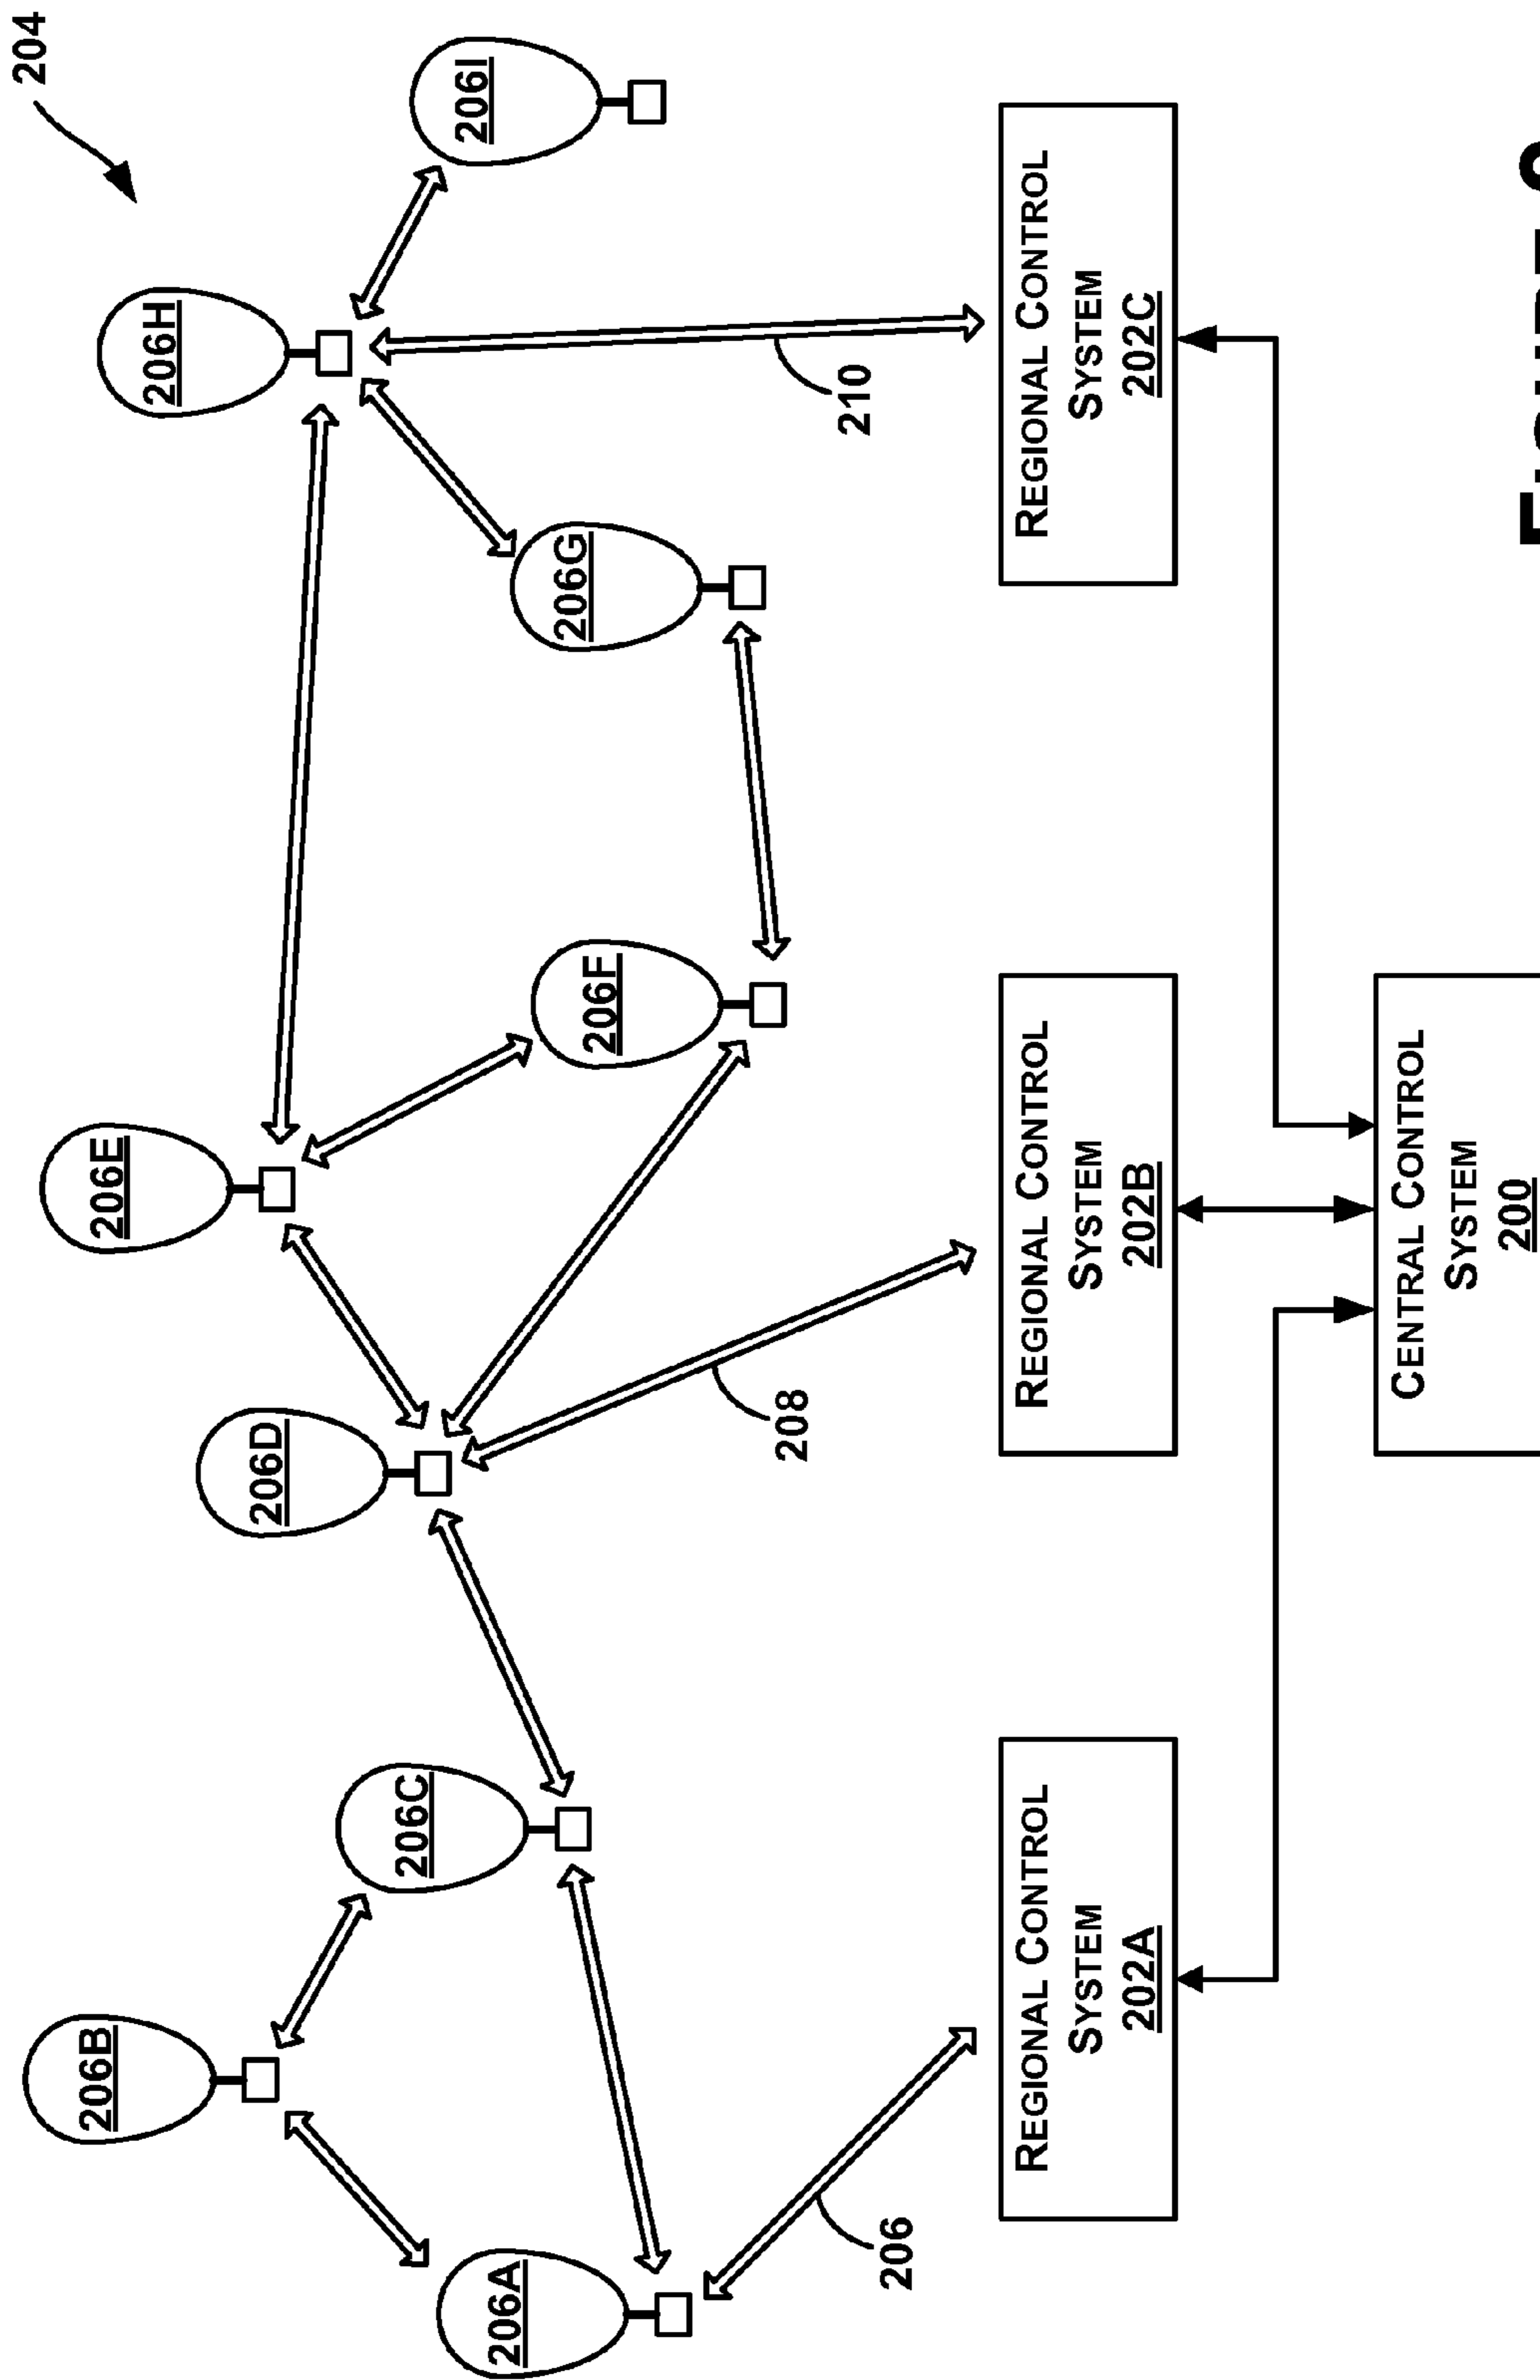
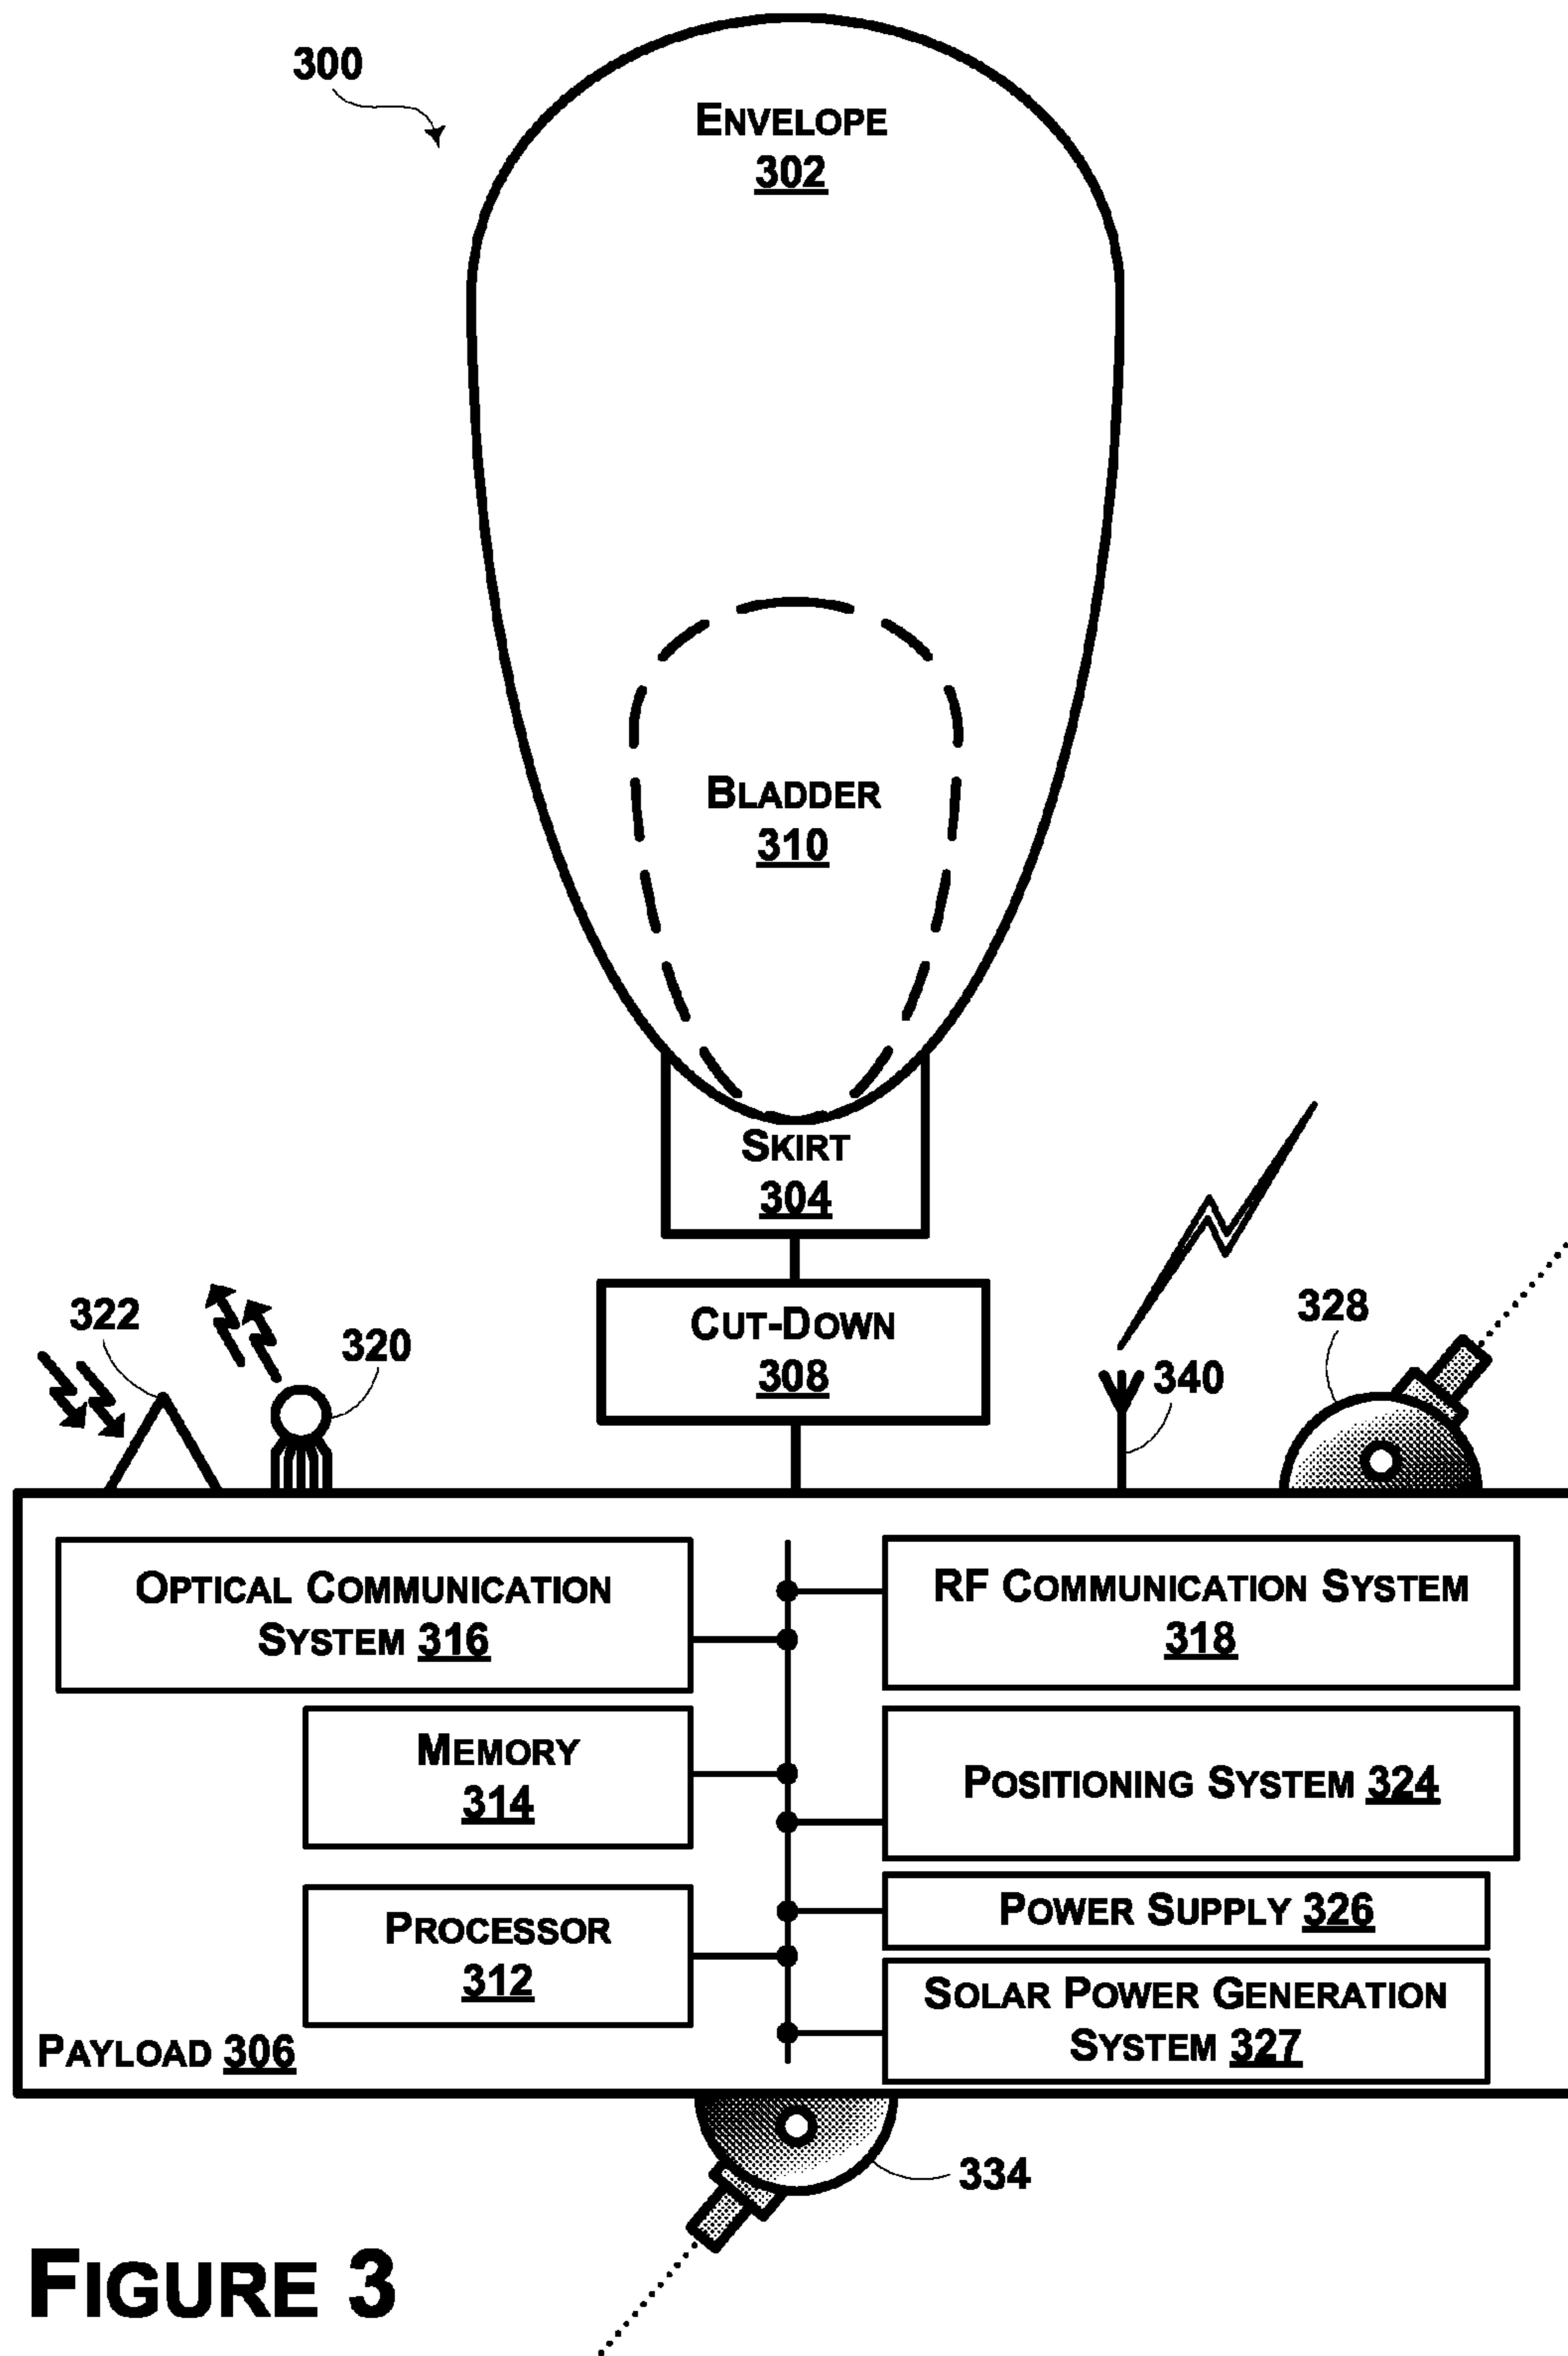


FIGURE 1



**FIGURE 2**



**FIGURE 3**



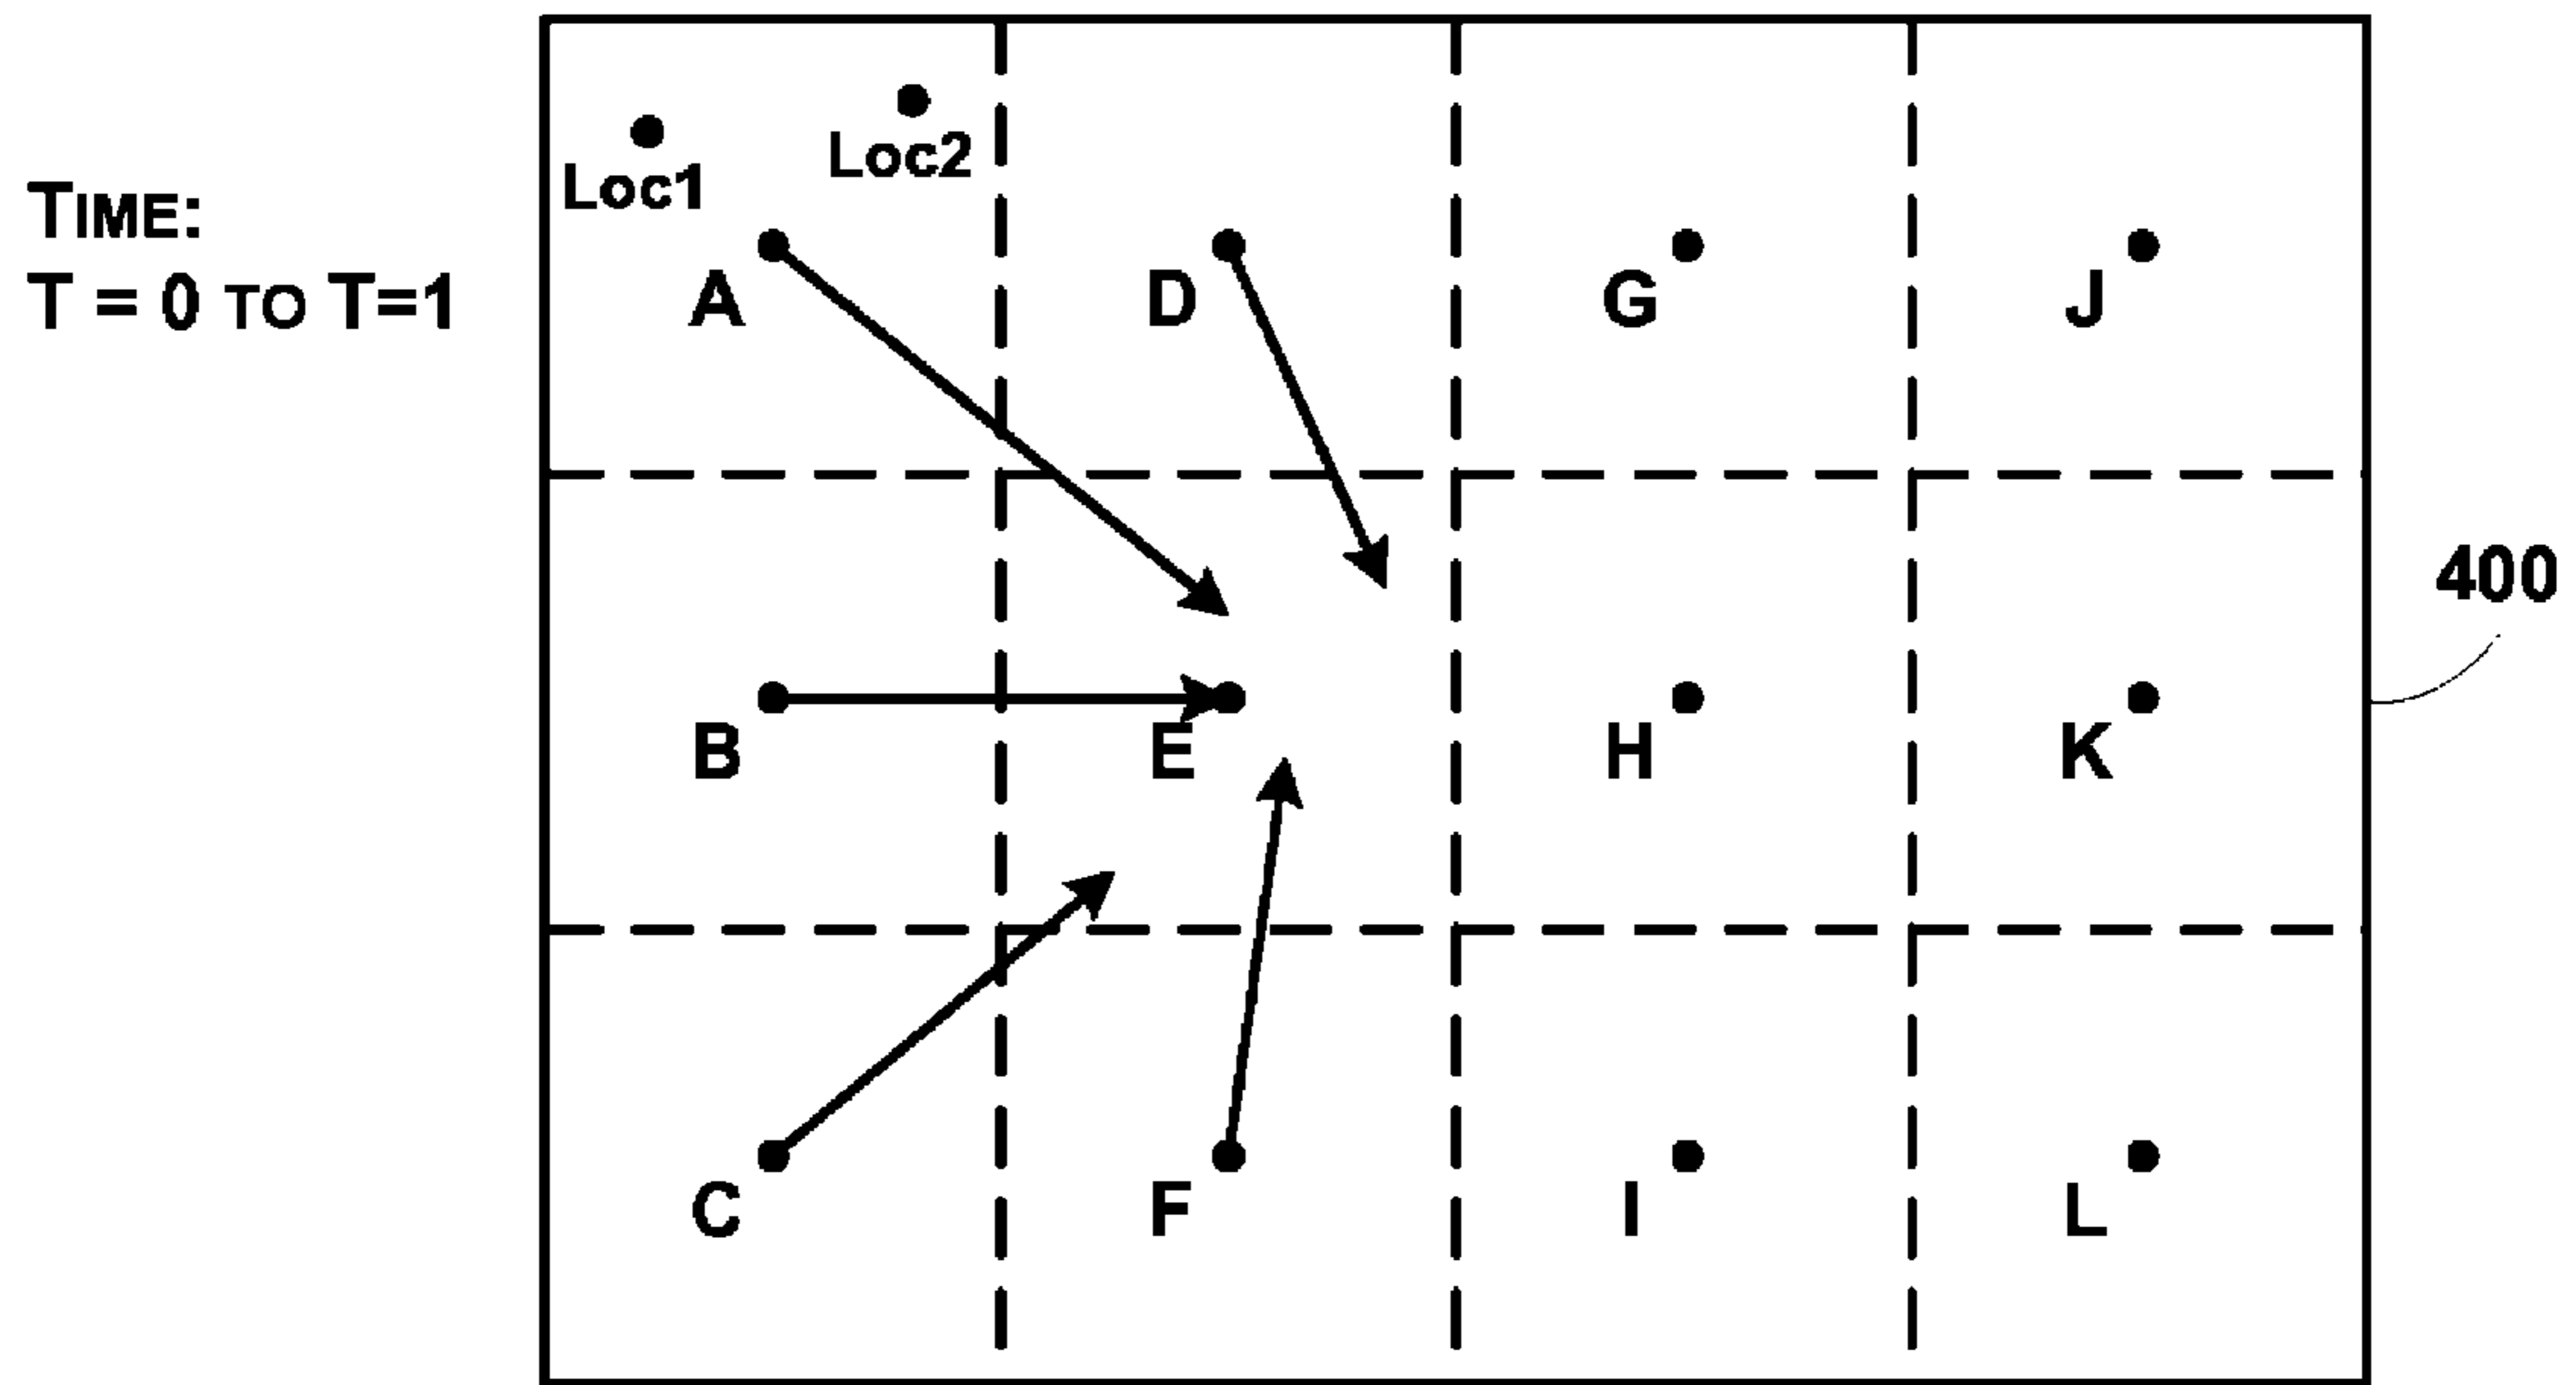


FIGURE 4

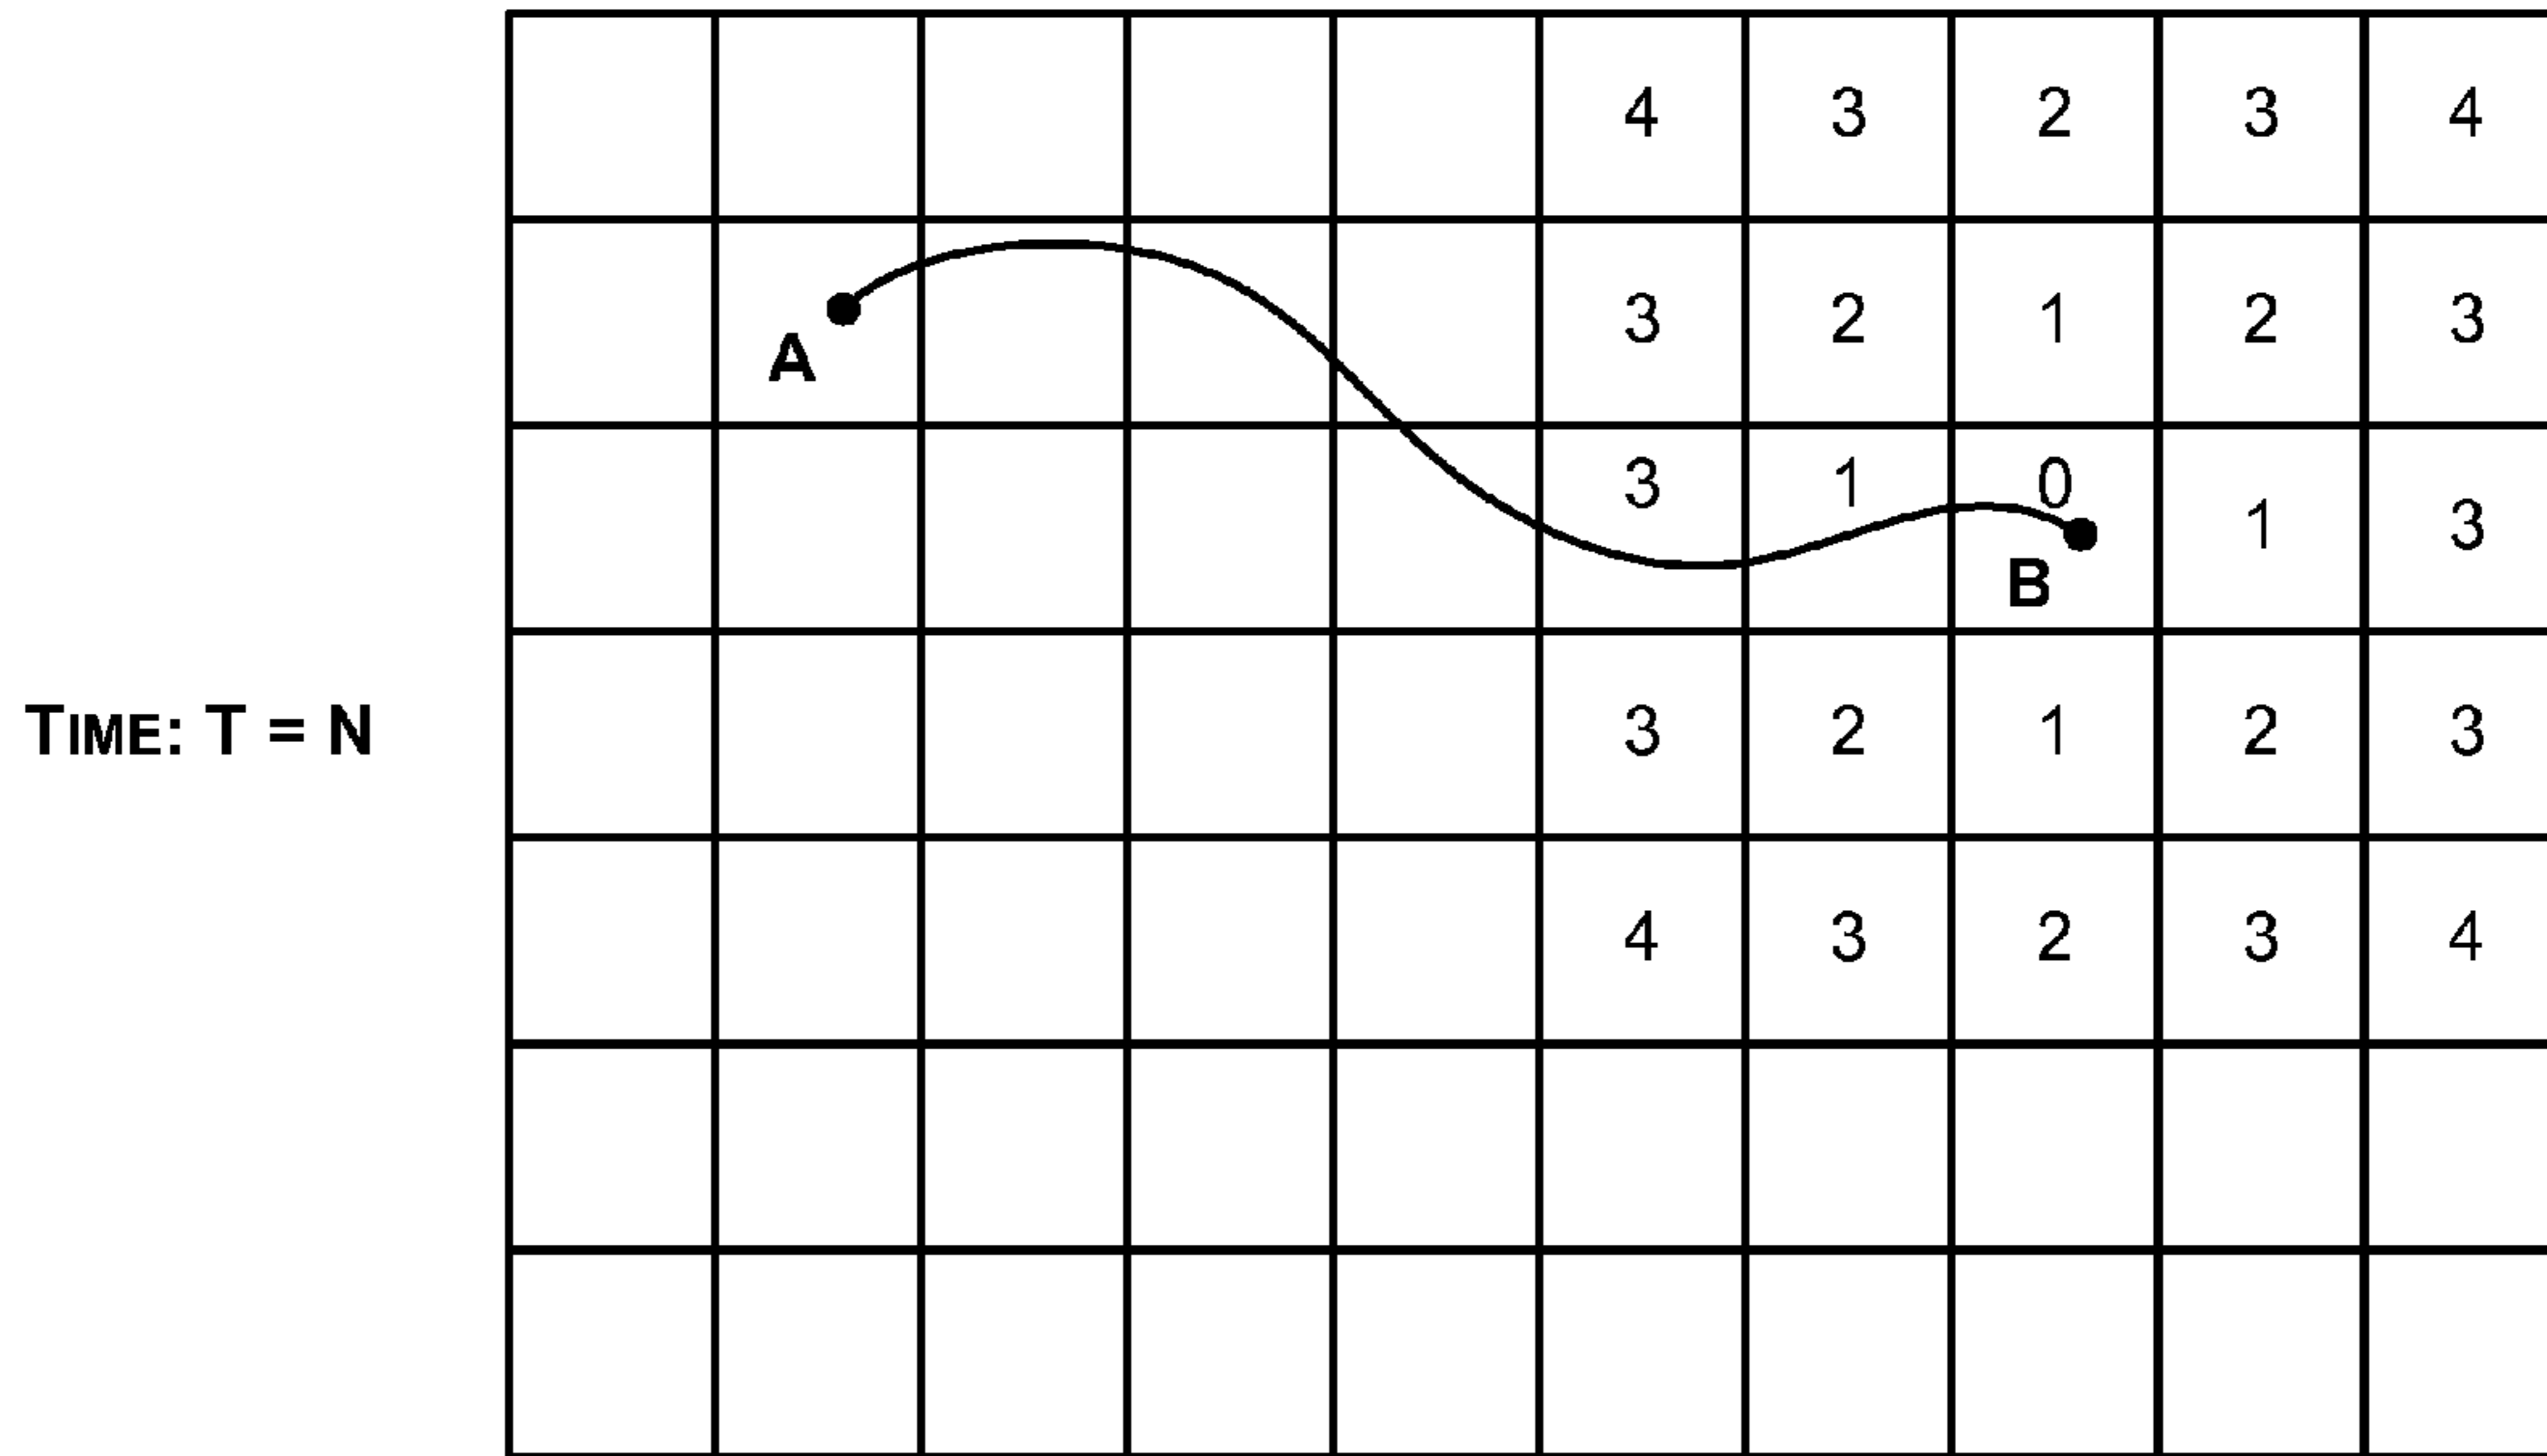


FIGURE 5

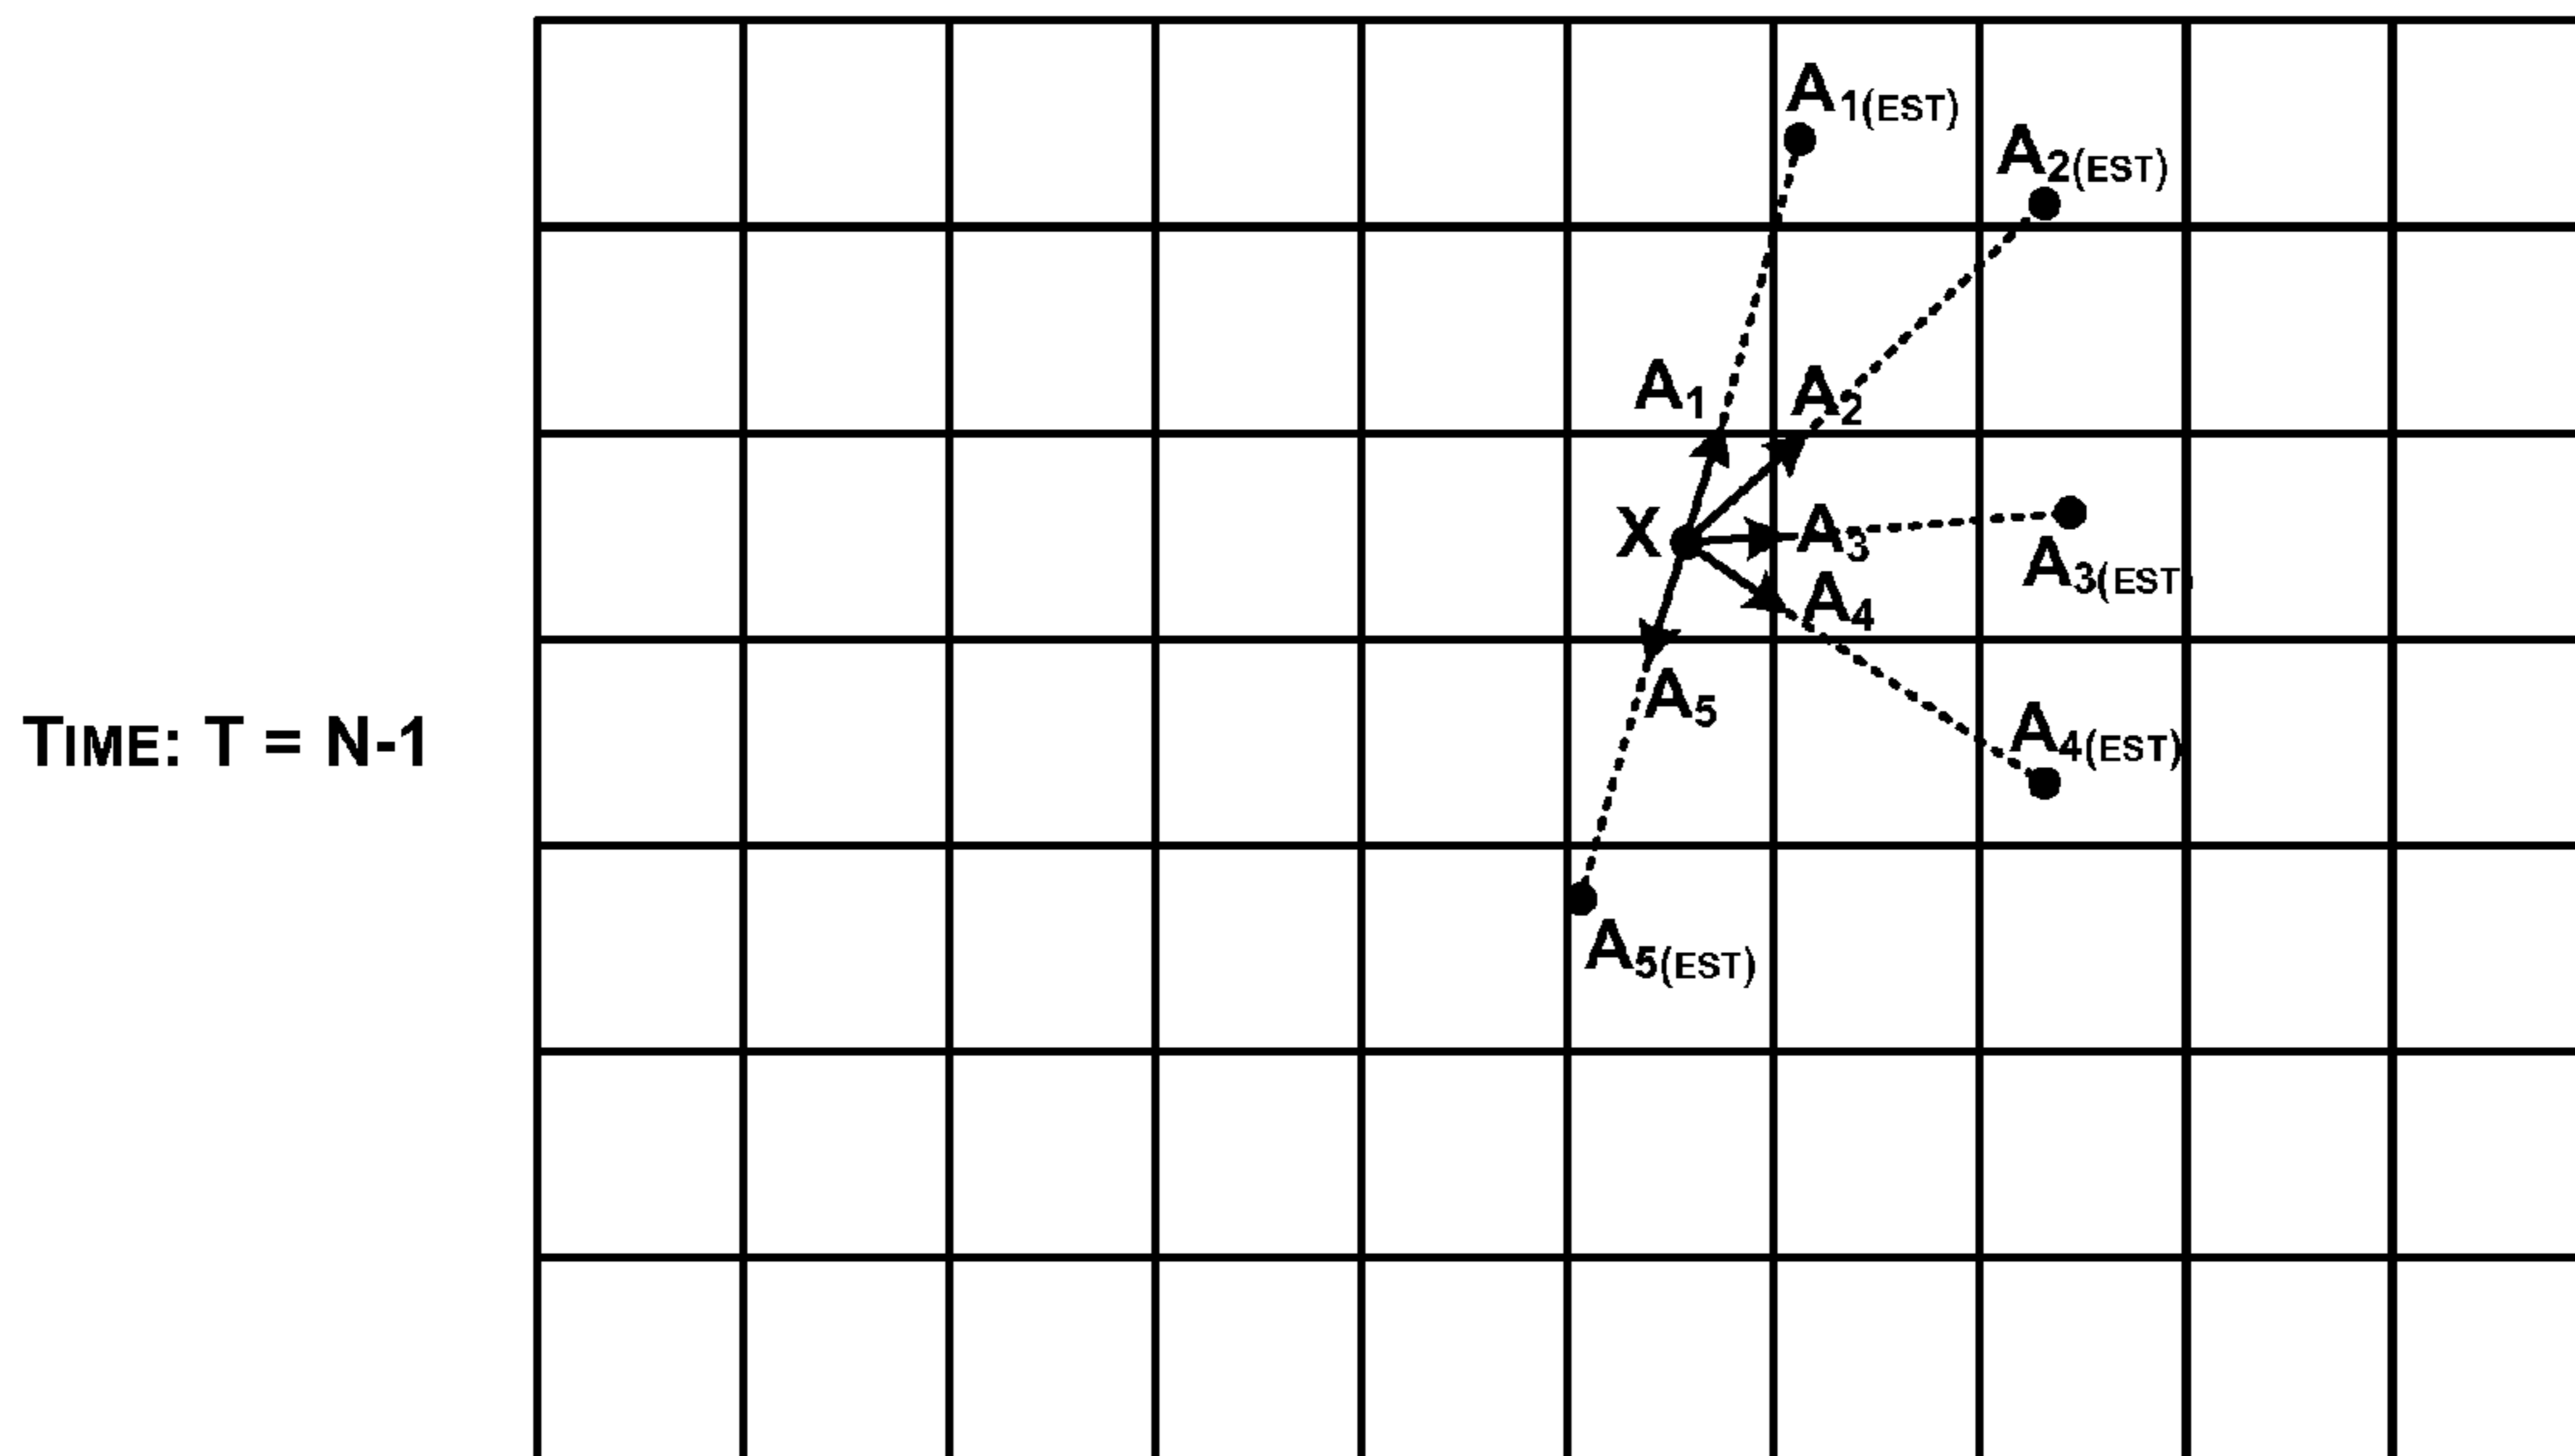


FIGURE 6

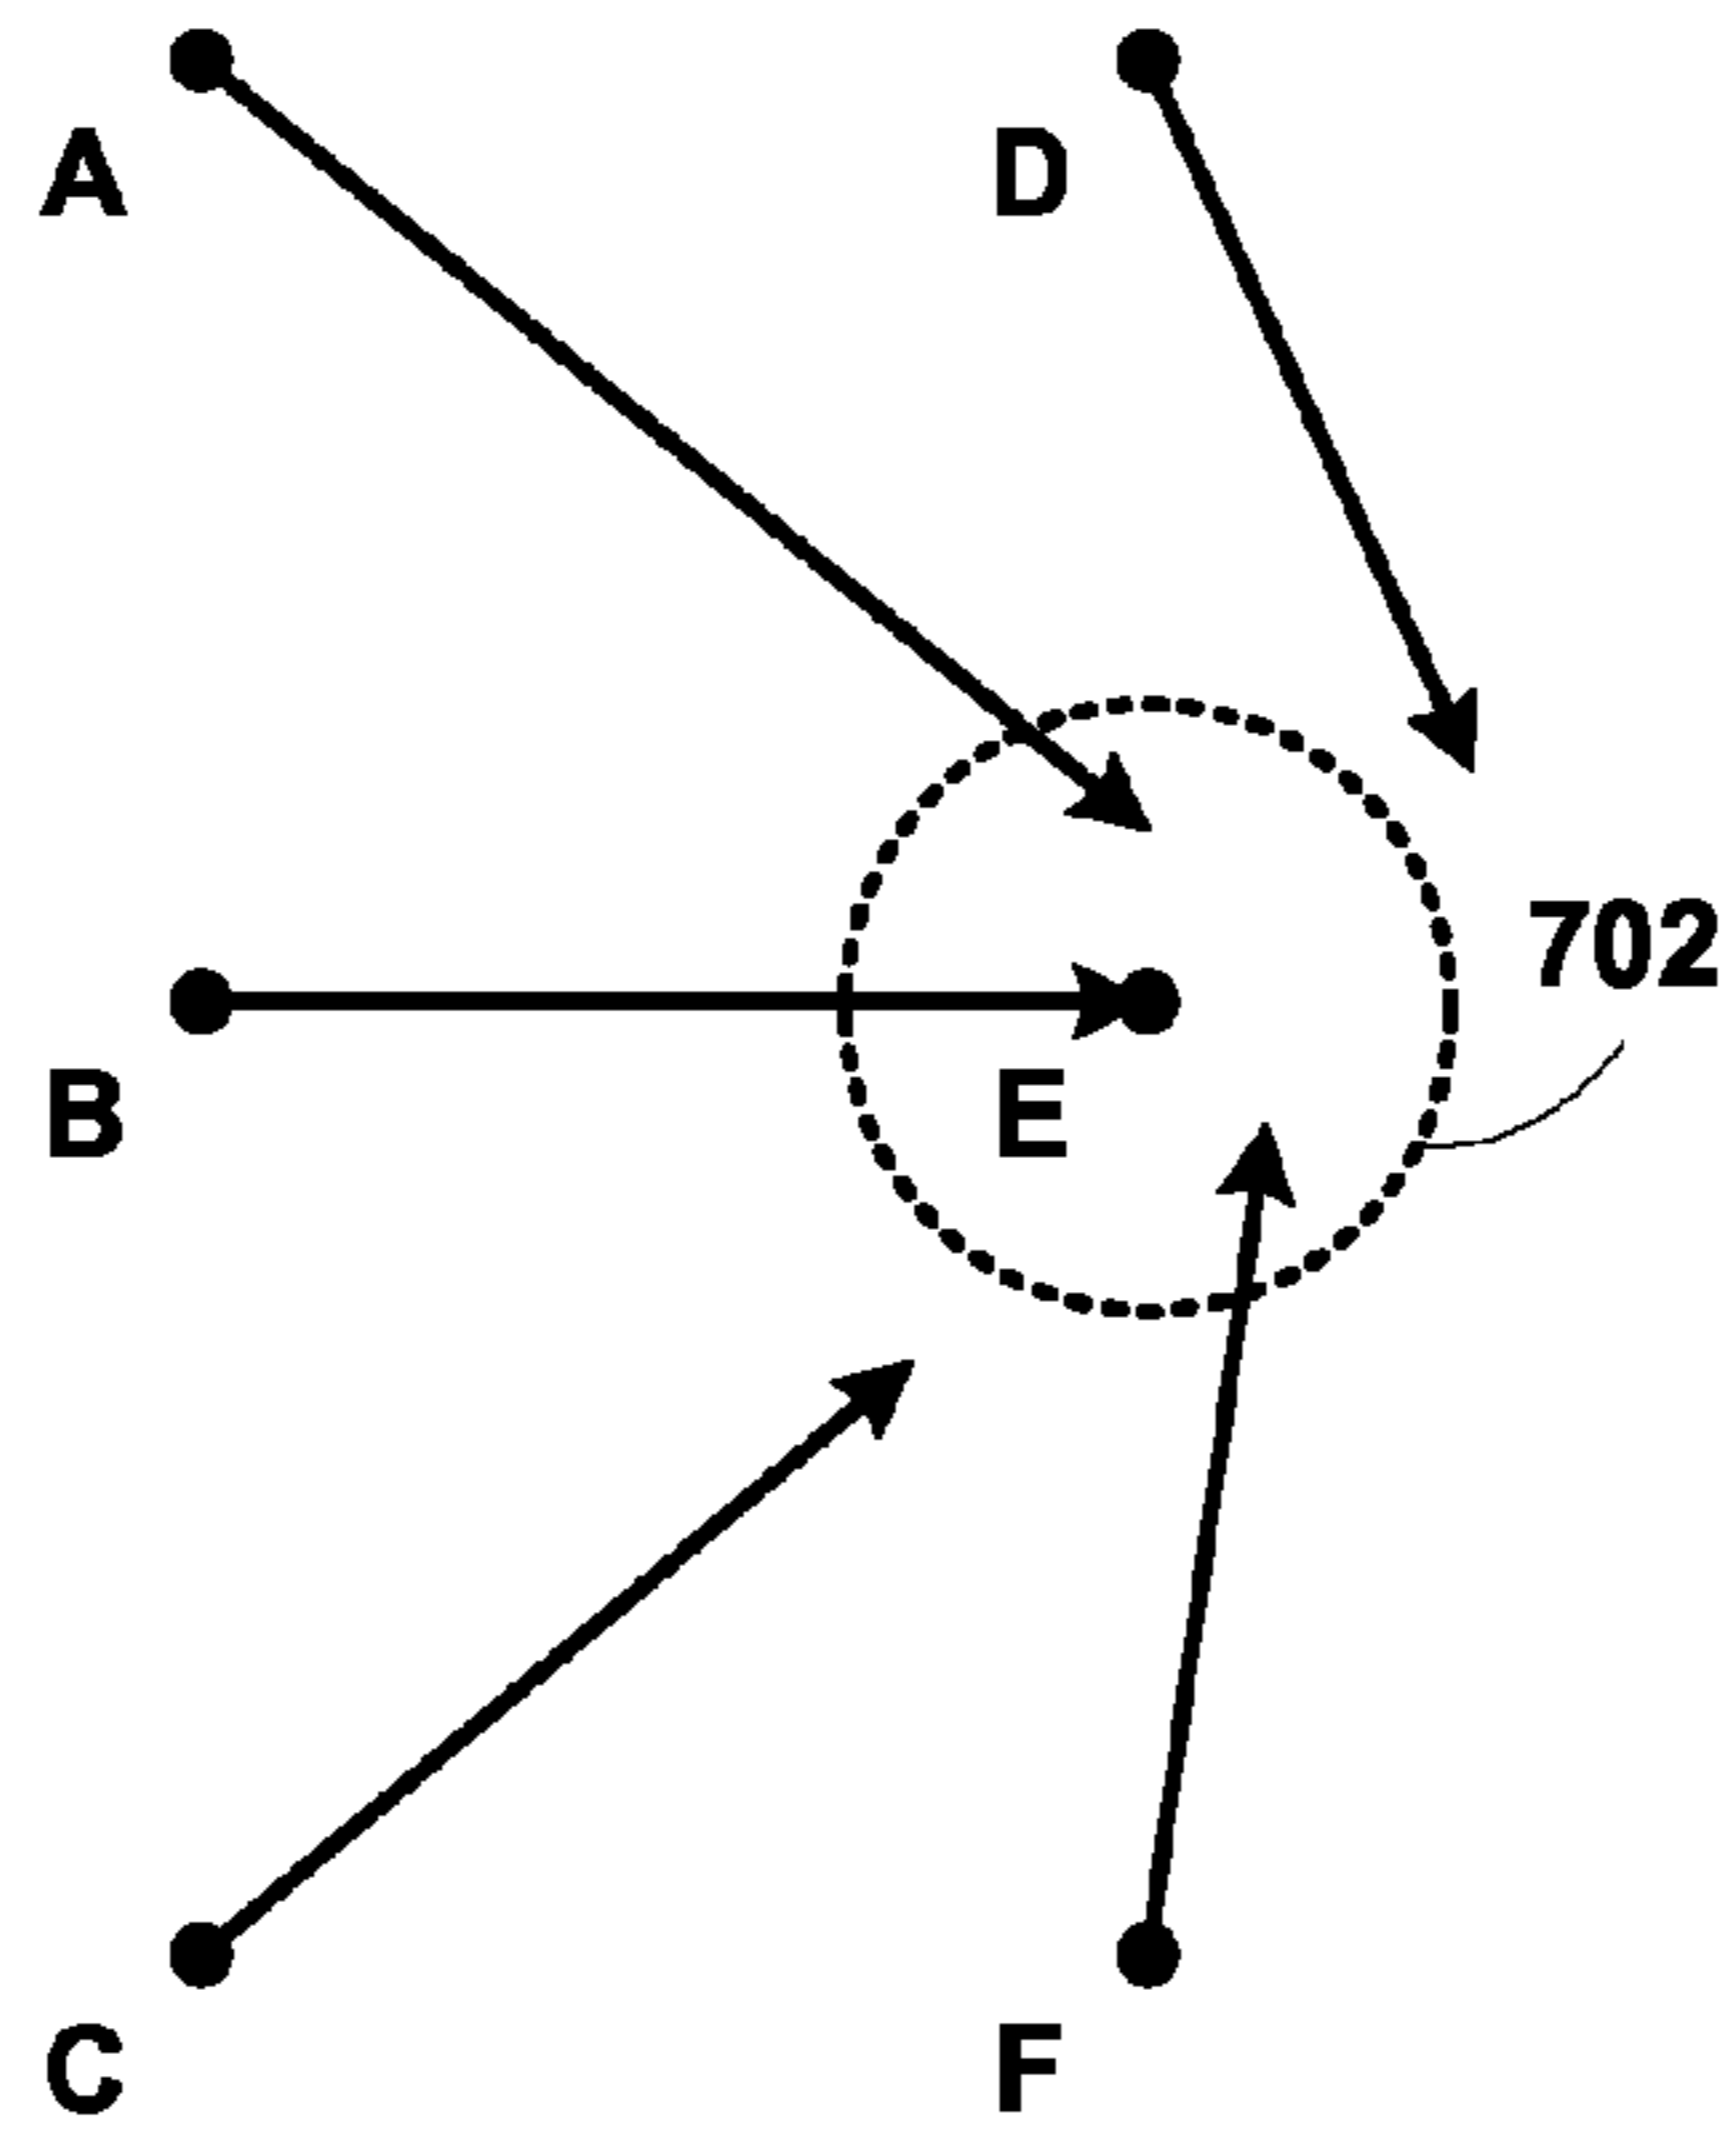


FIGURE 7A

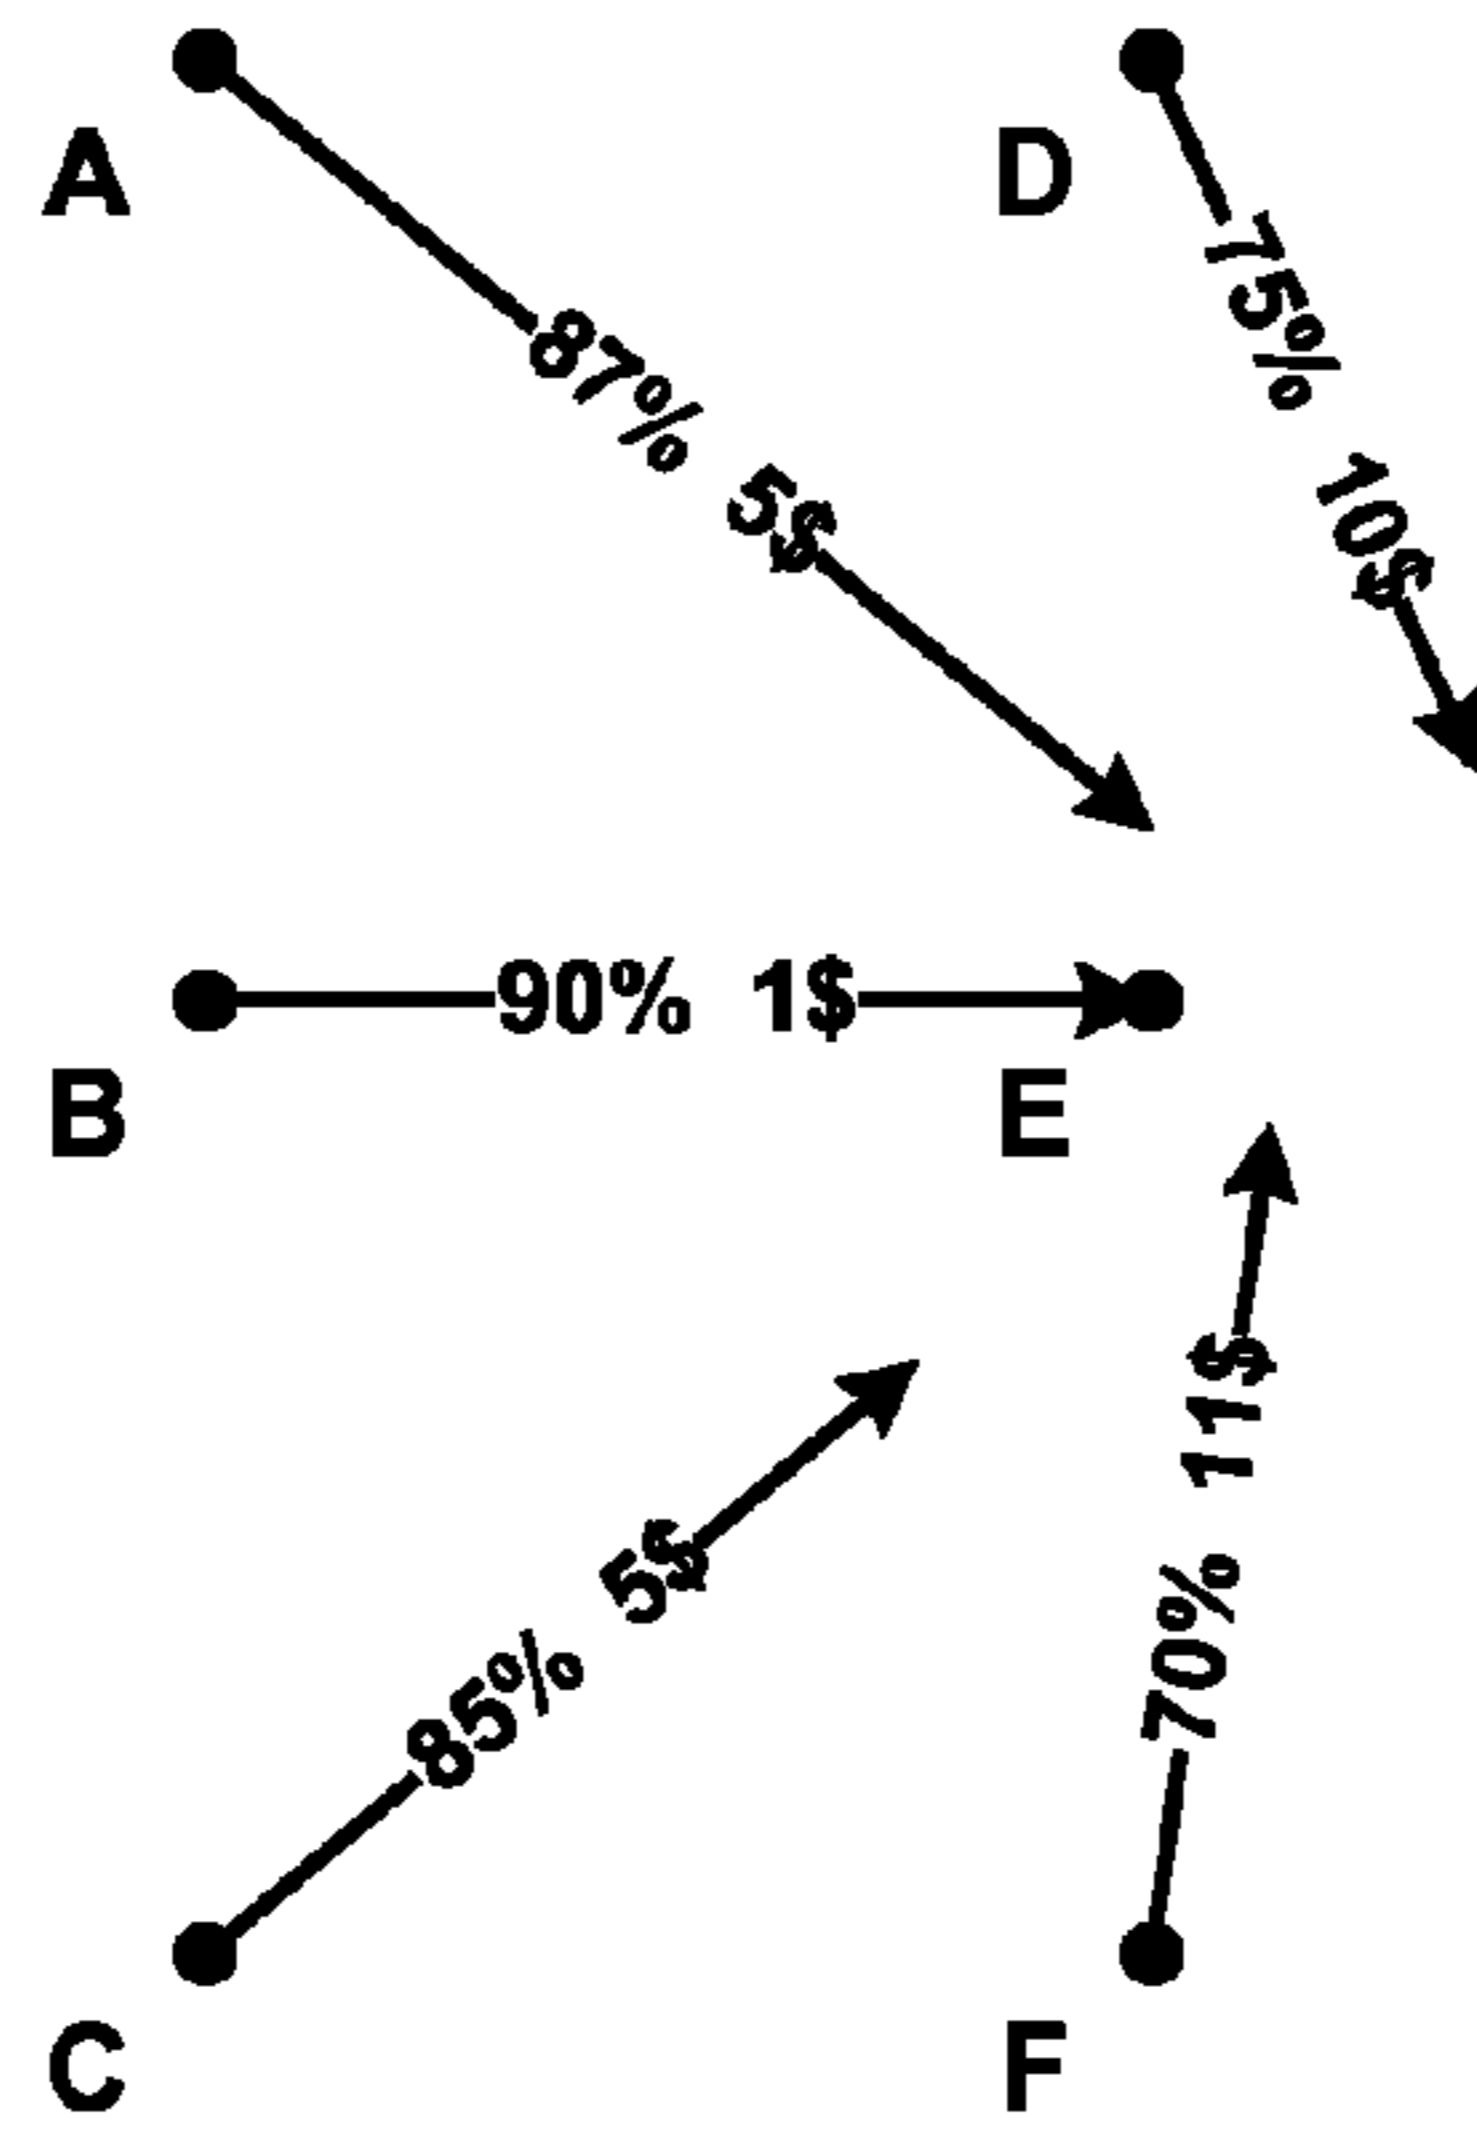


FIGURE 7B



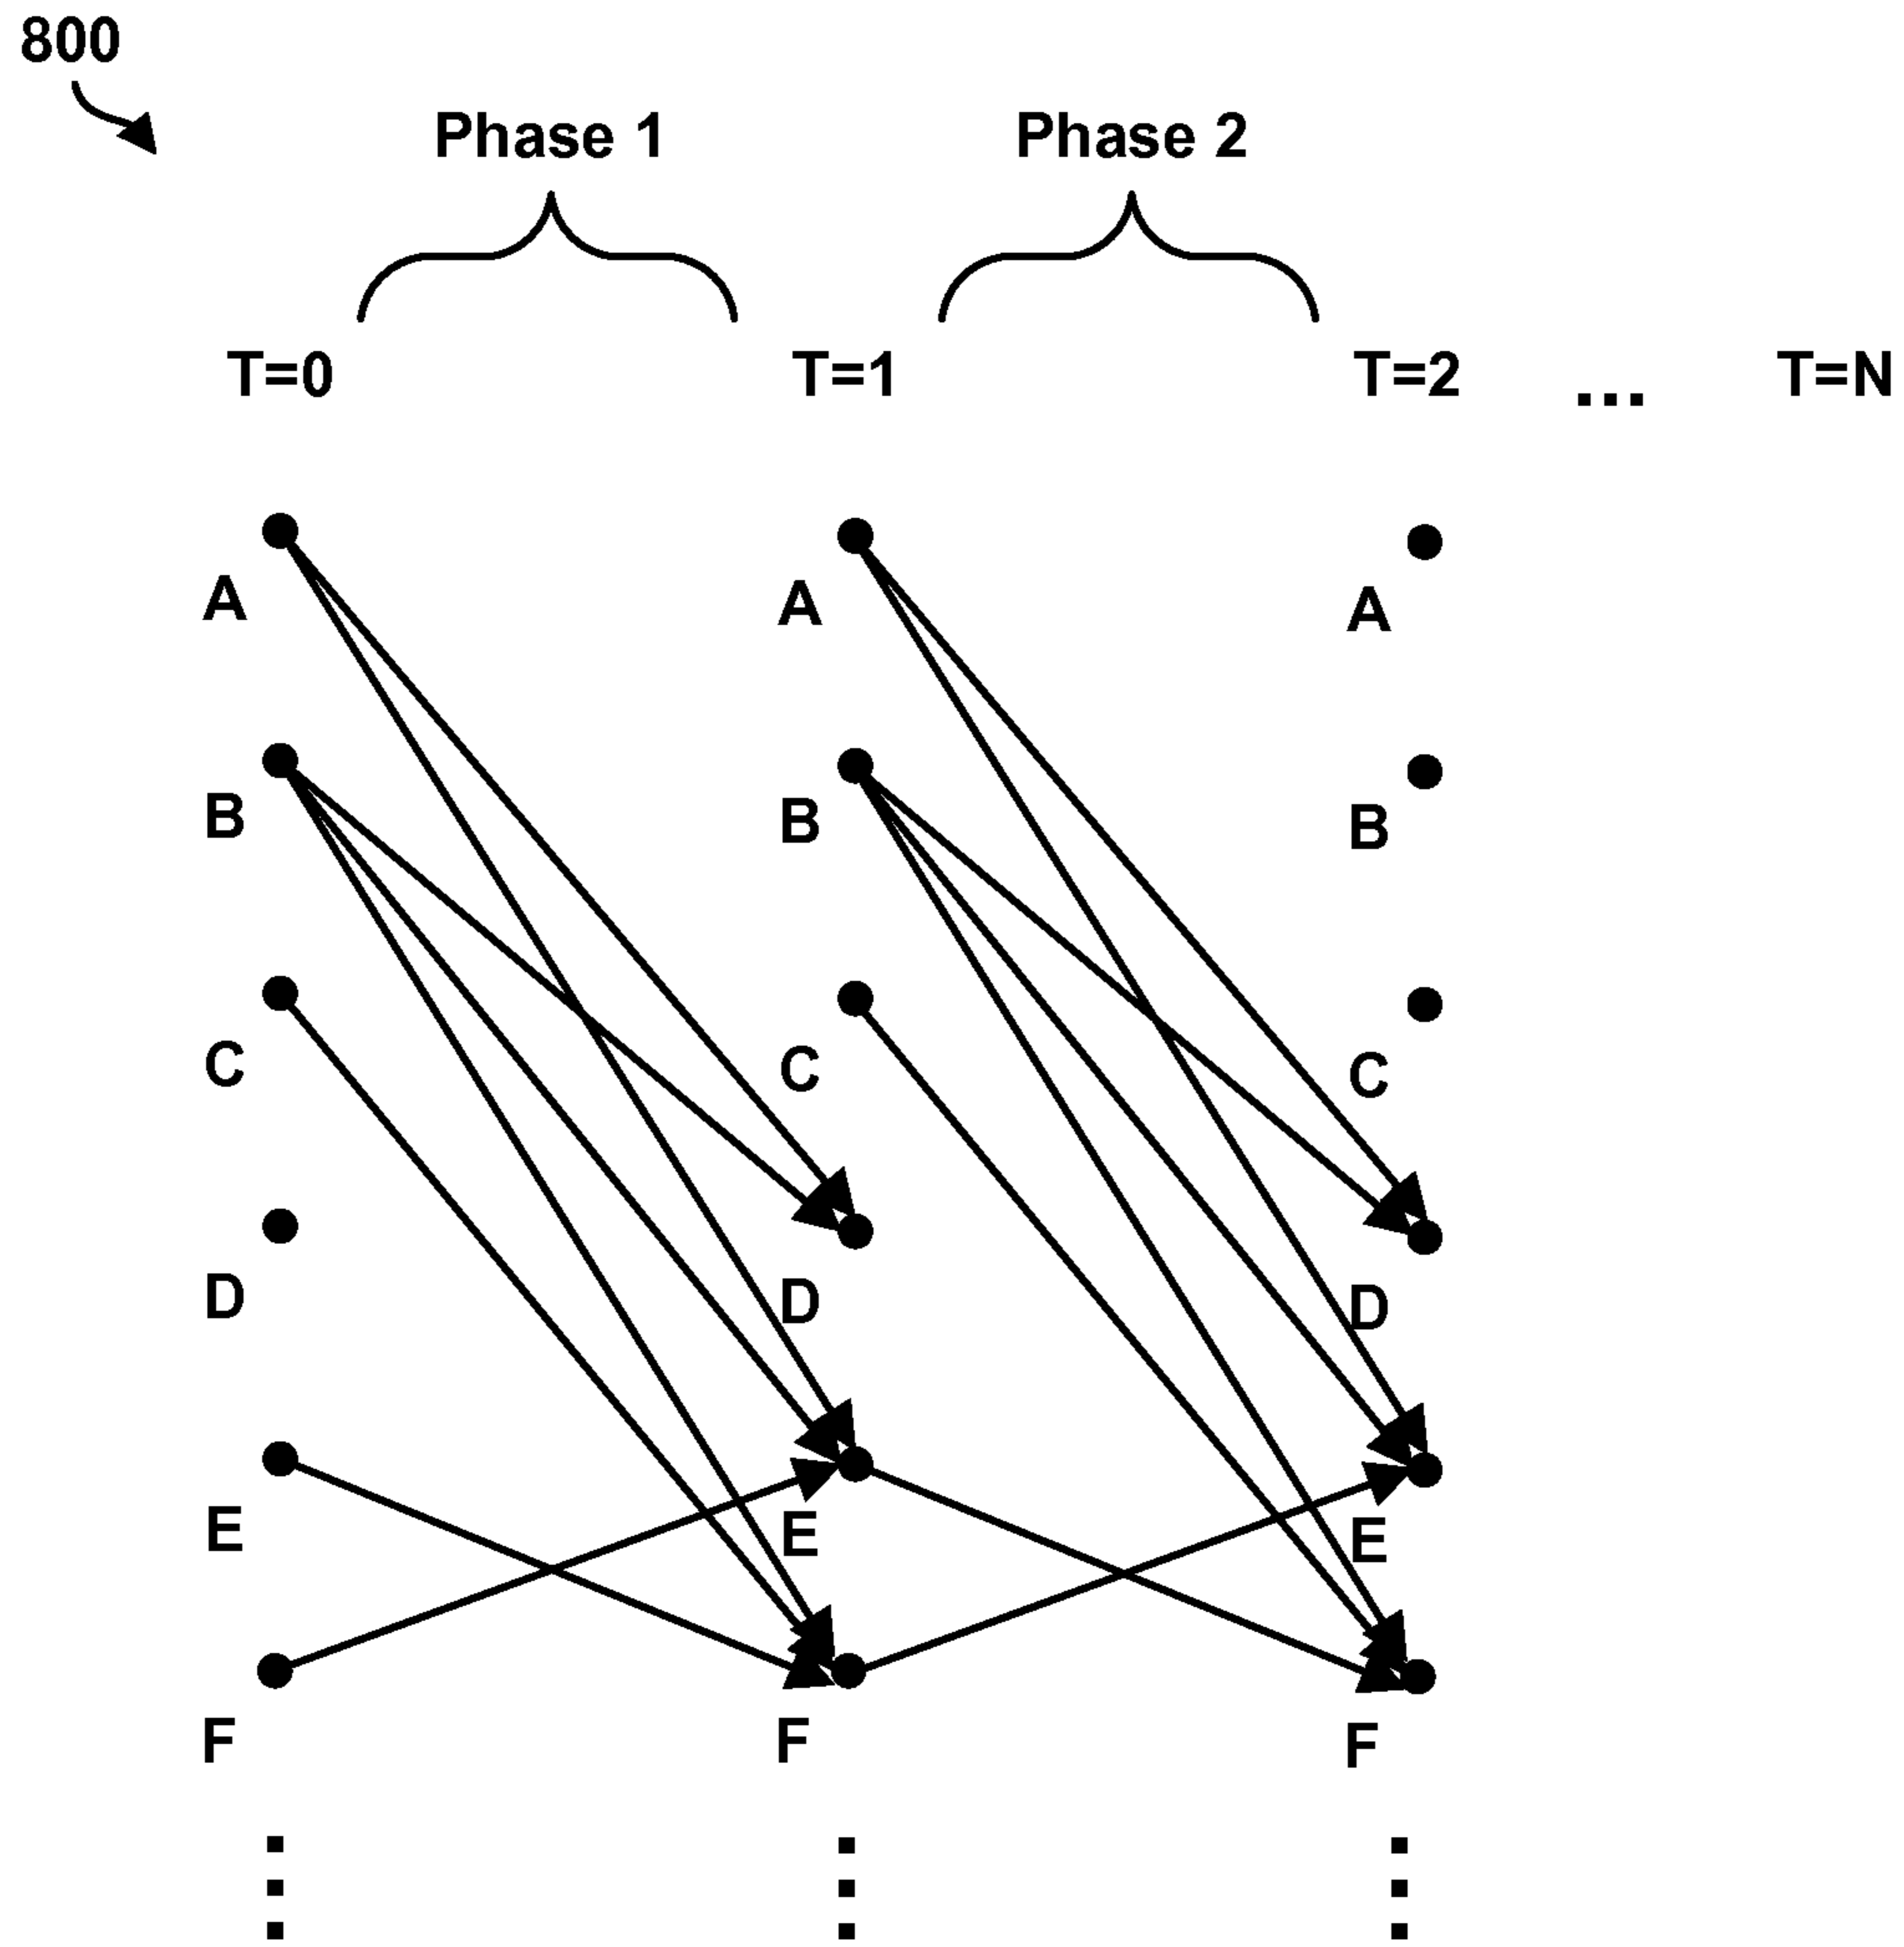
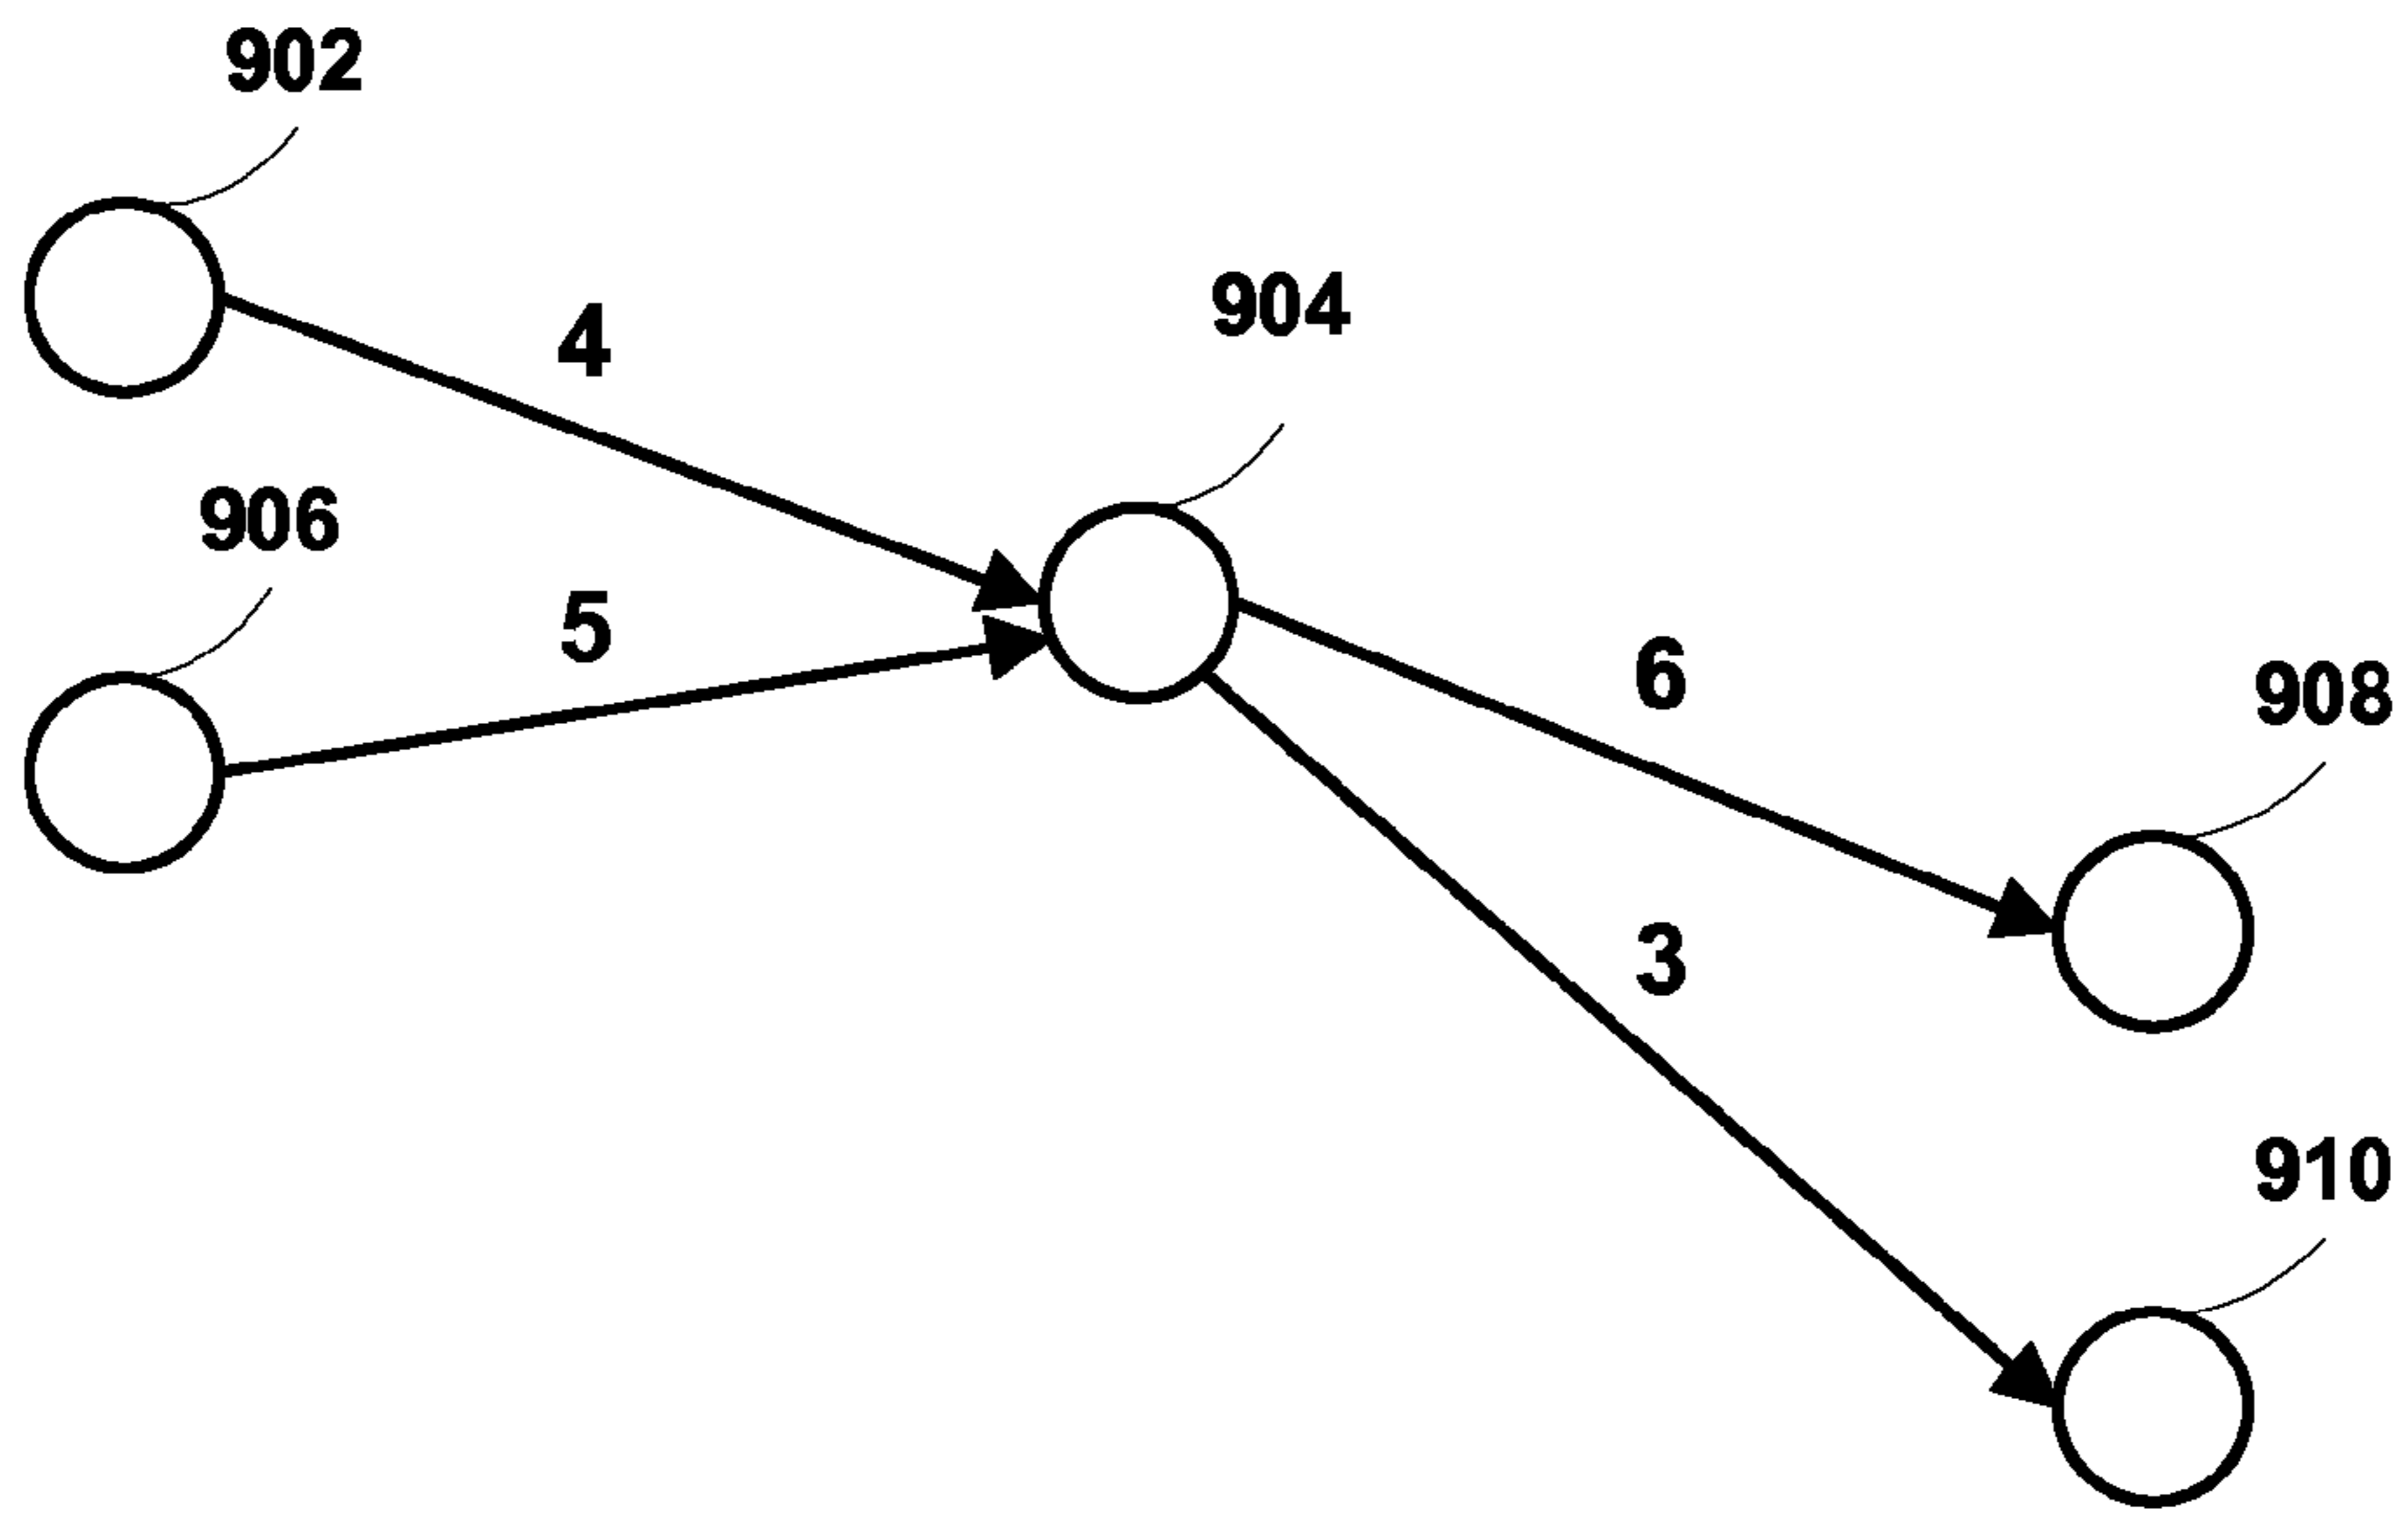
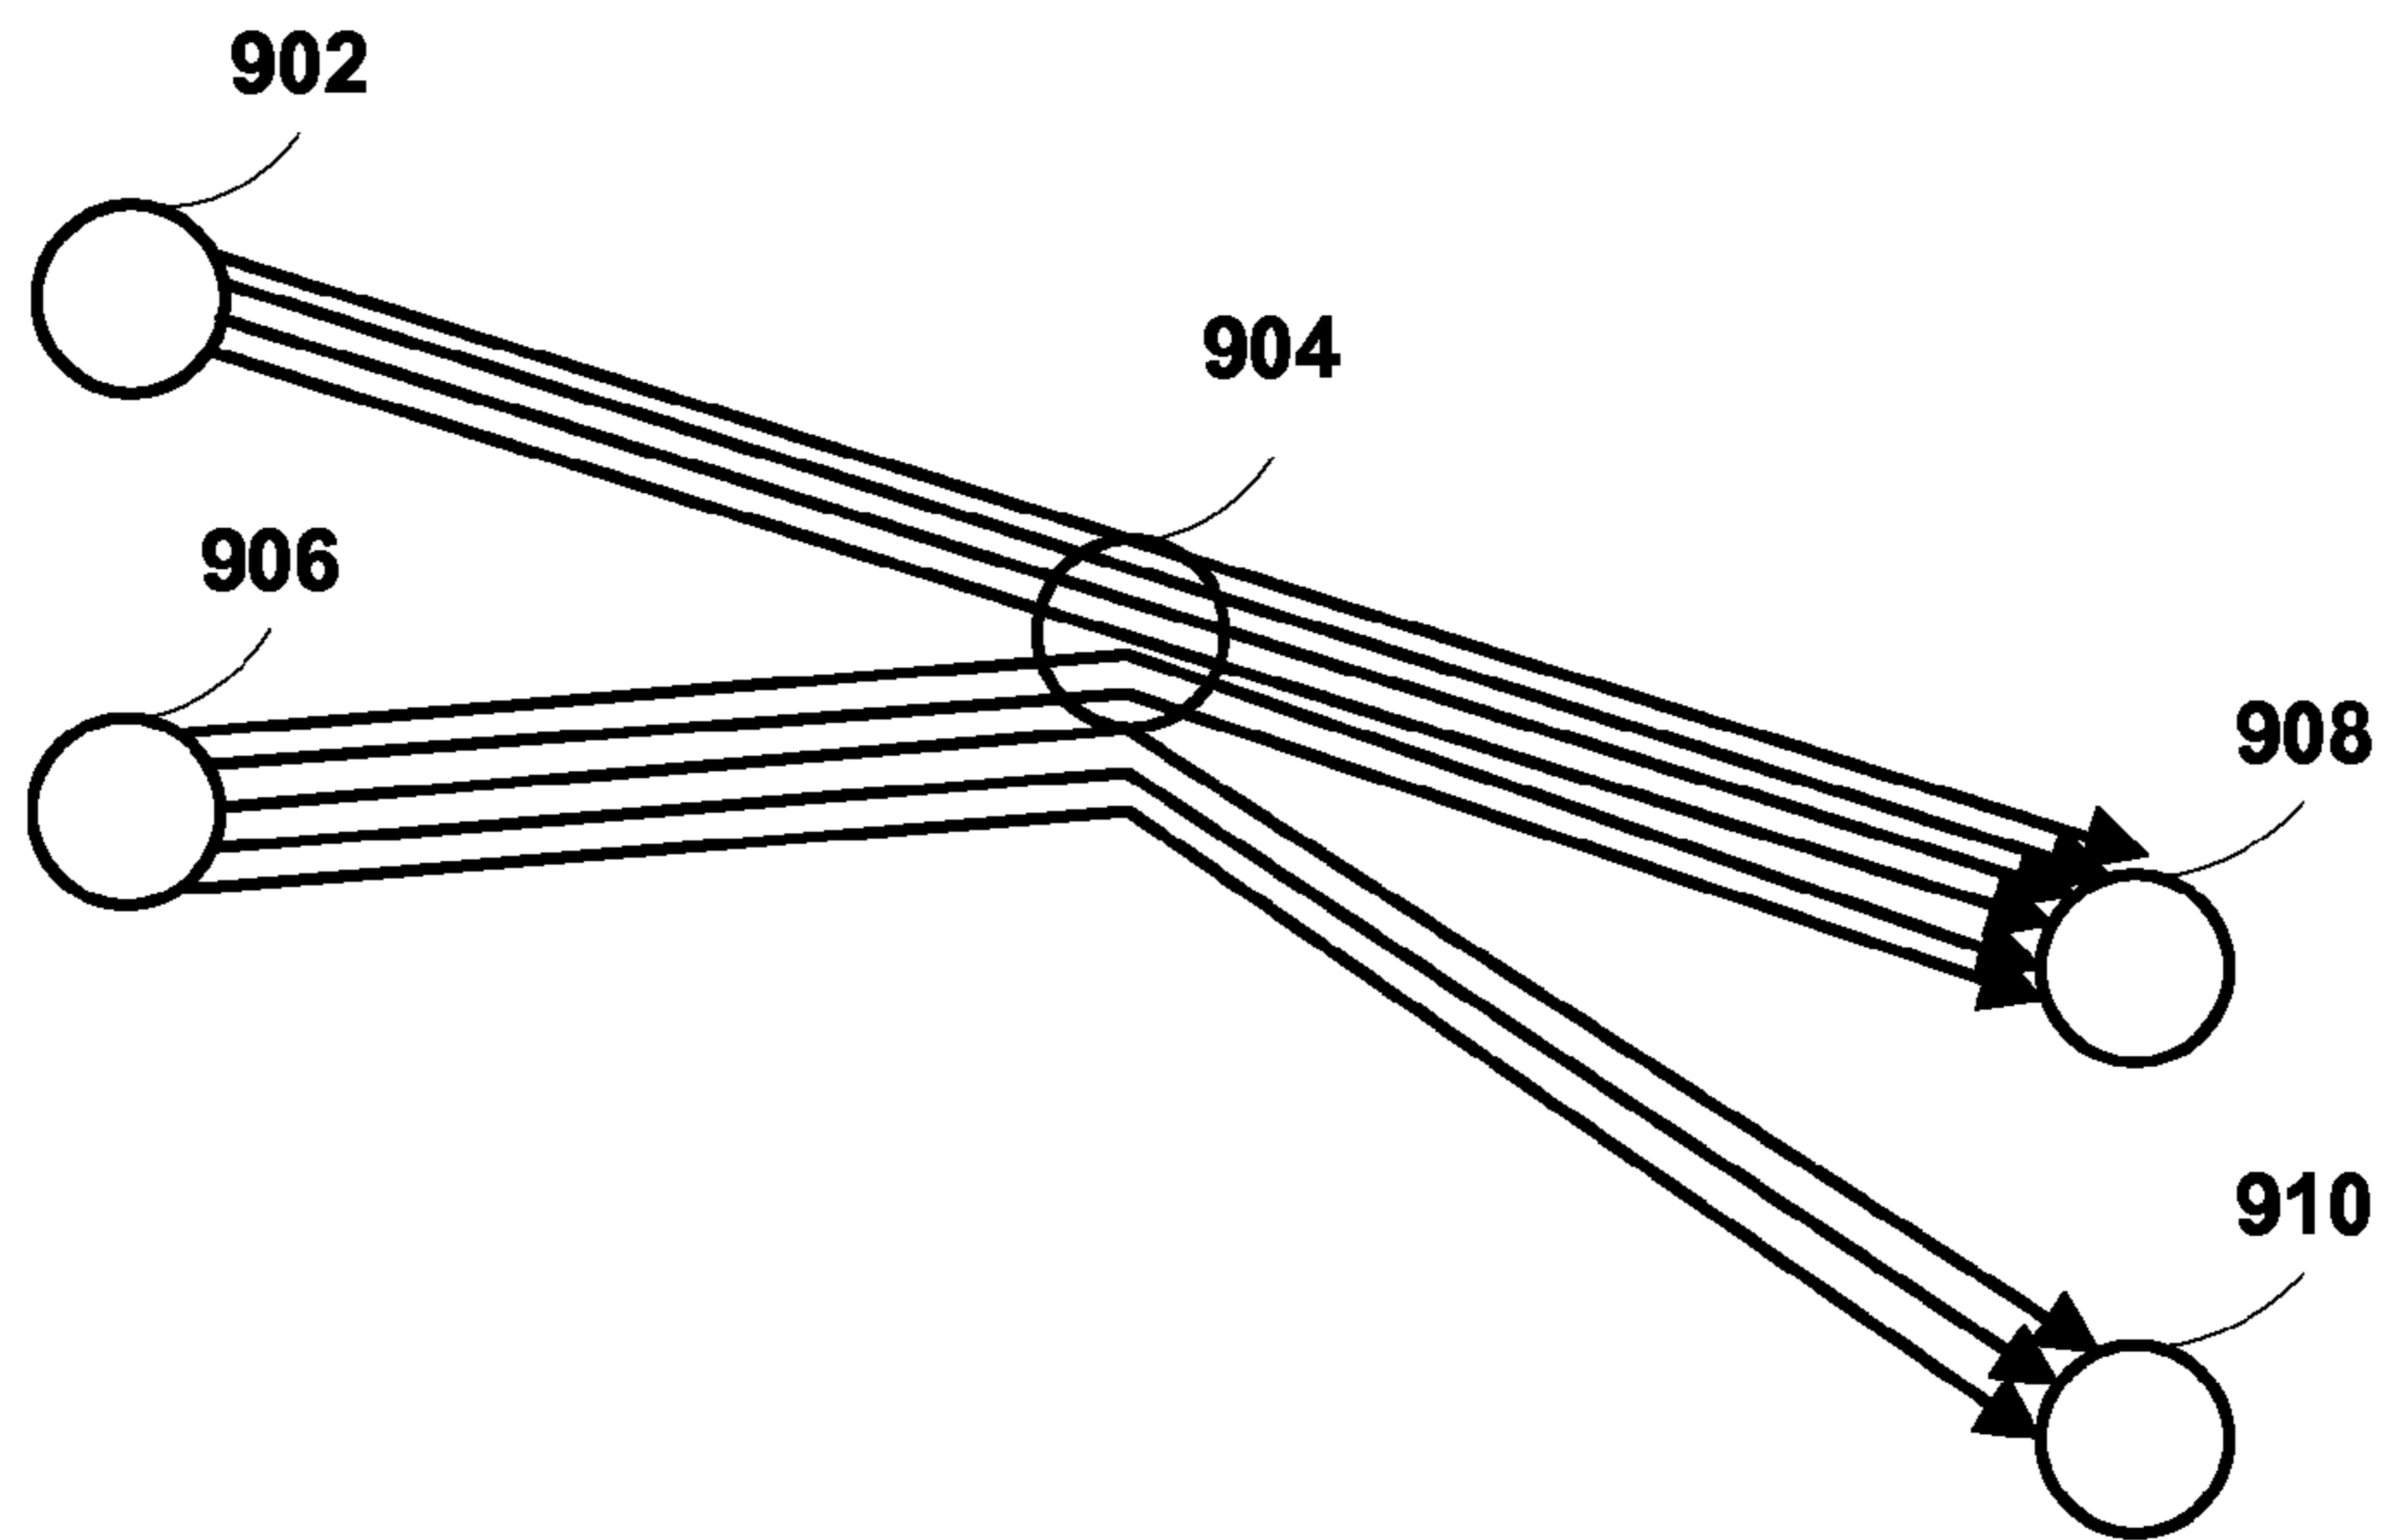


FIGURE 8



**FIGURE 9A**



**FIGURE 9B**

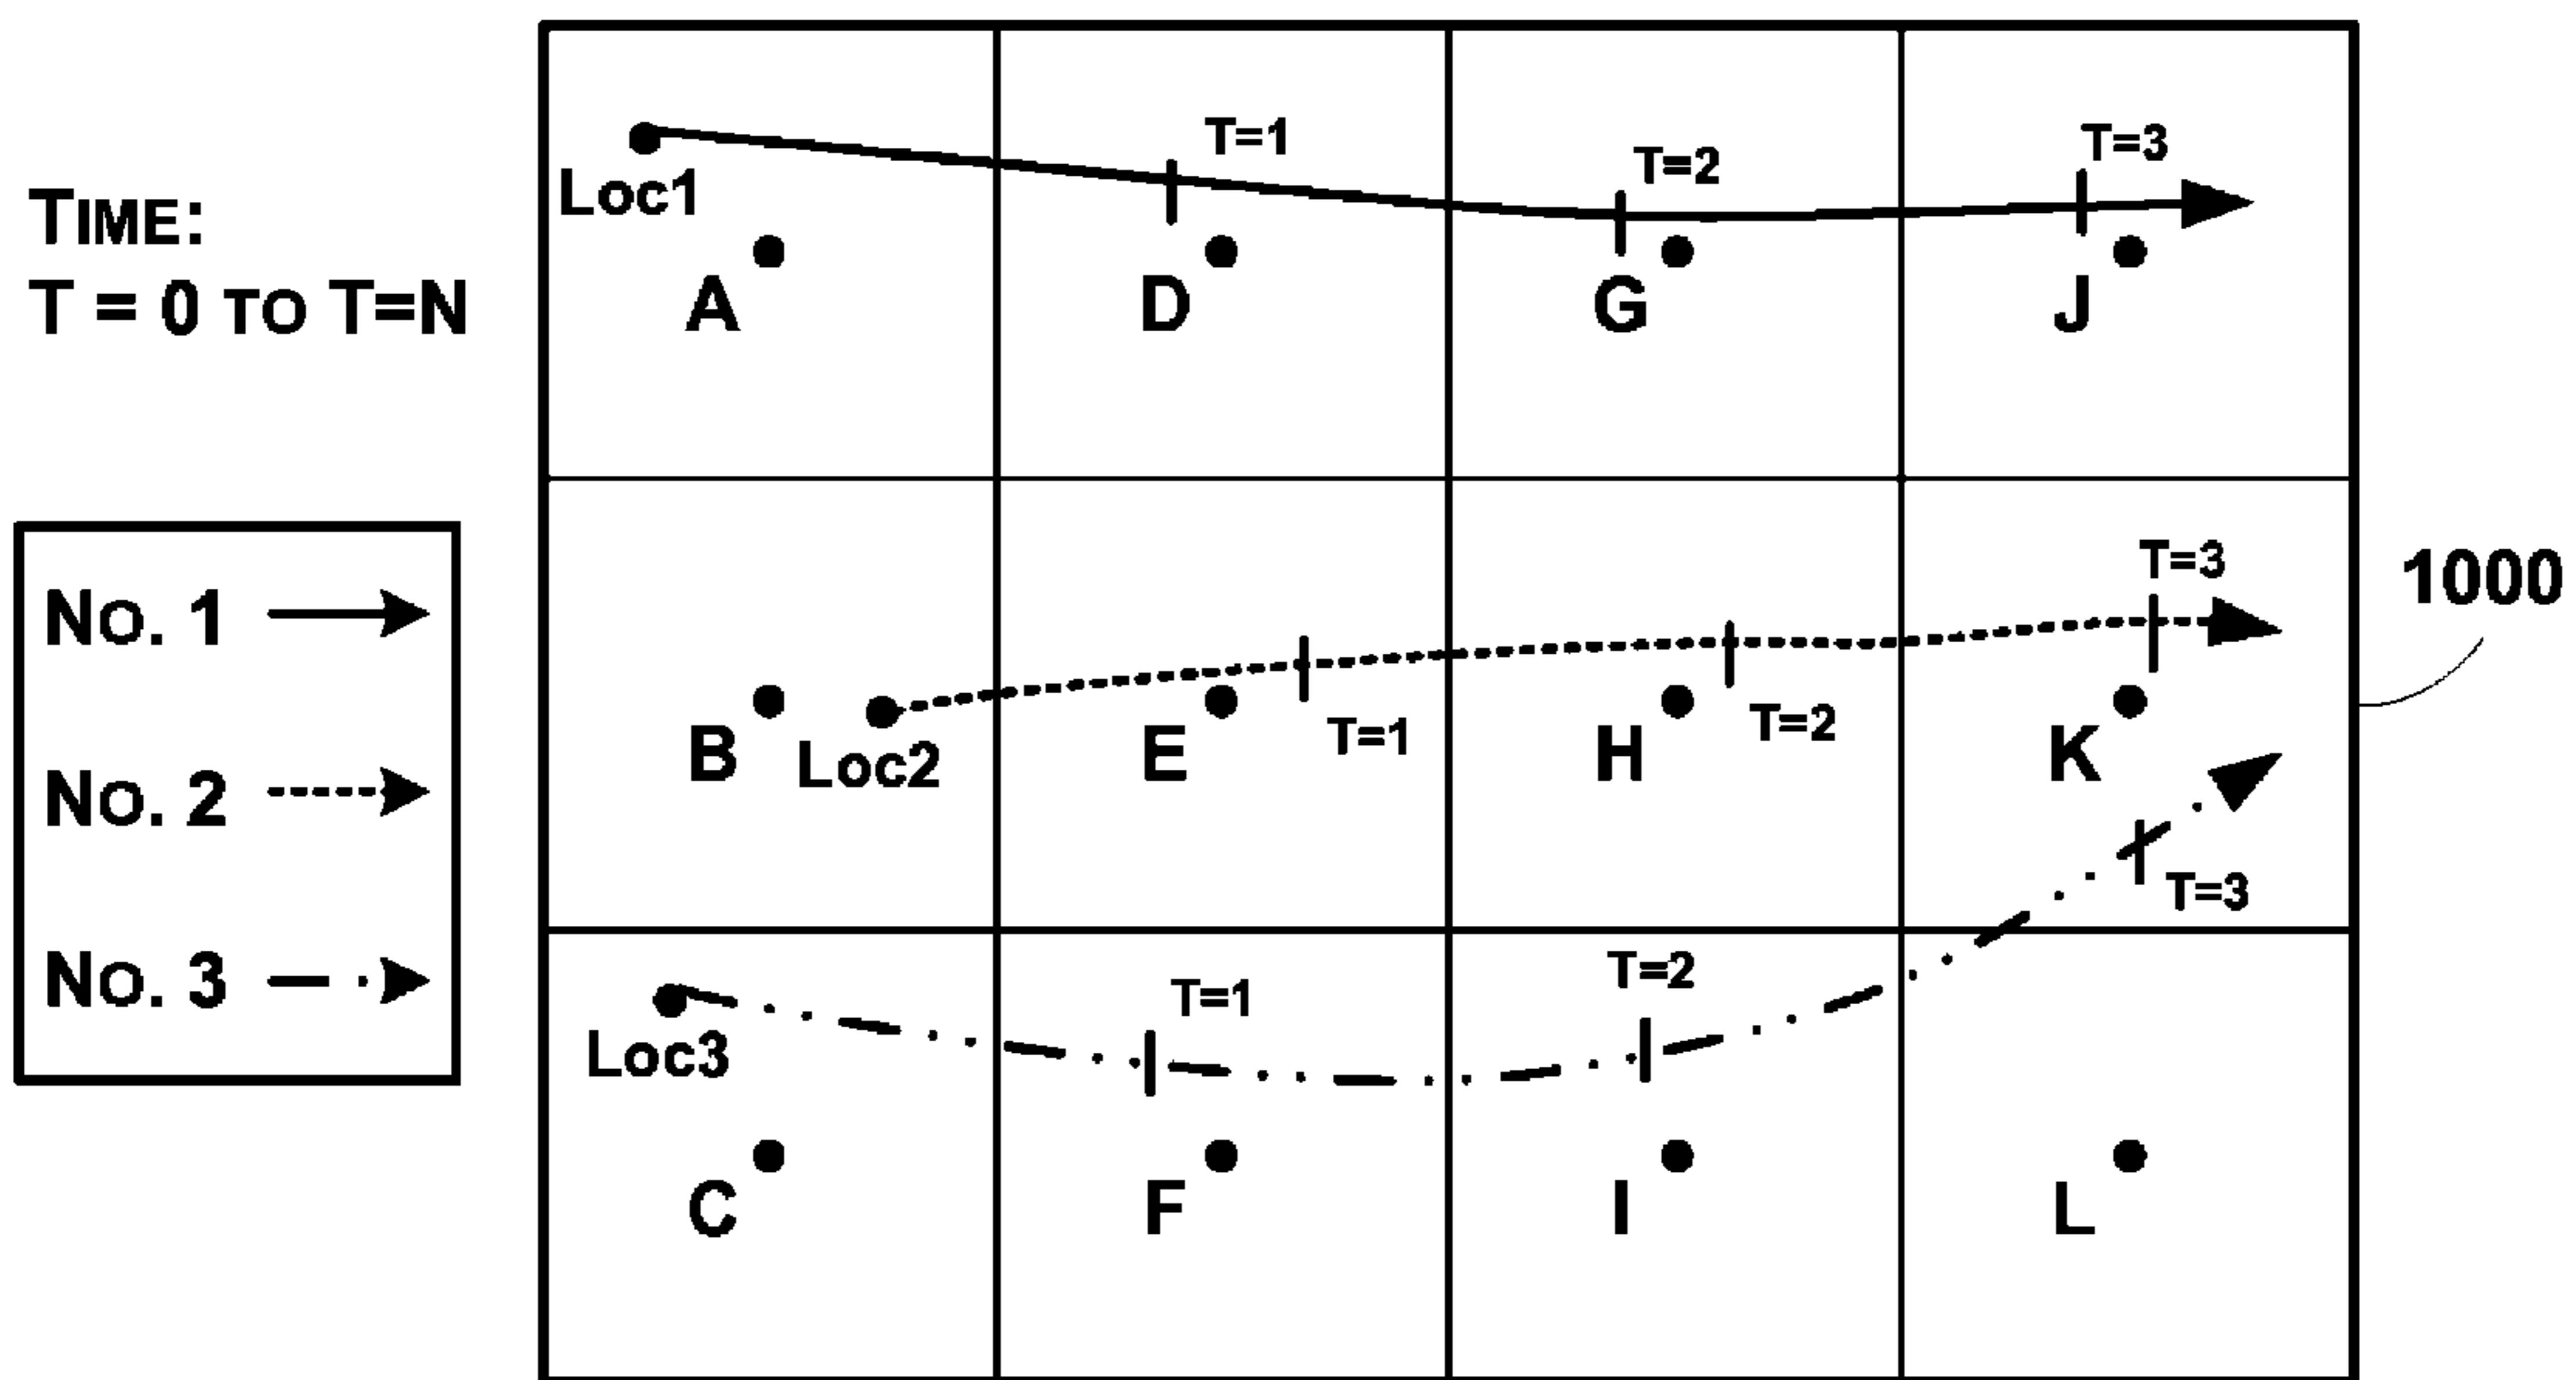
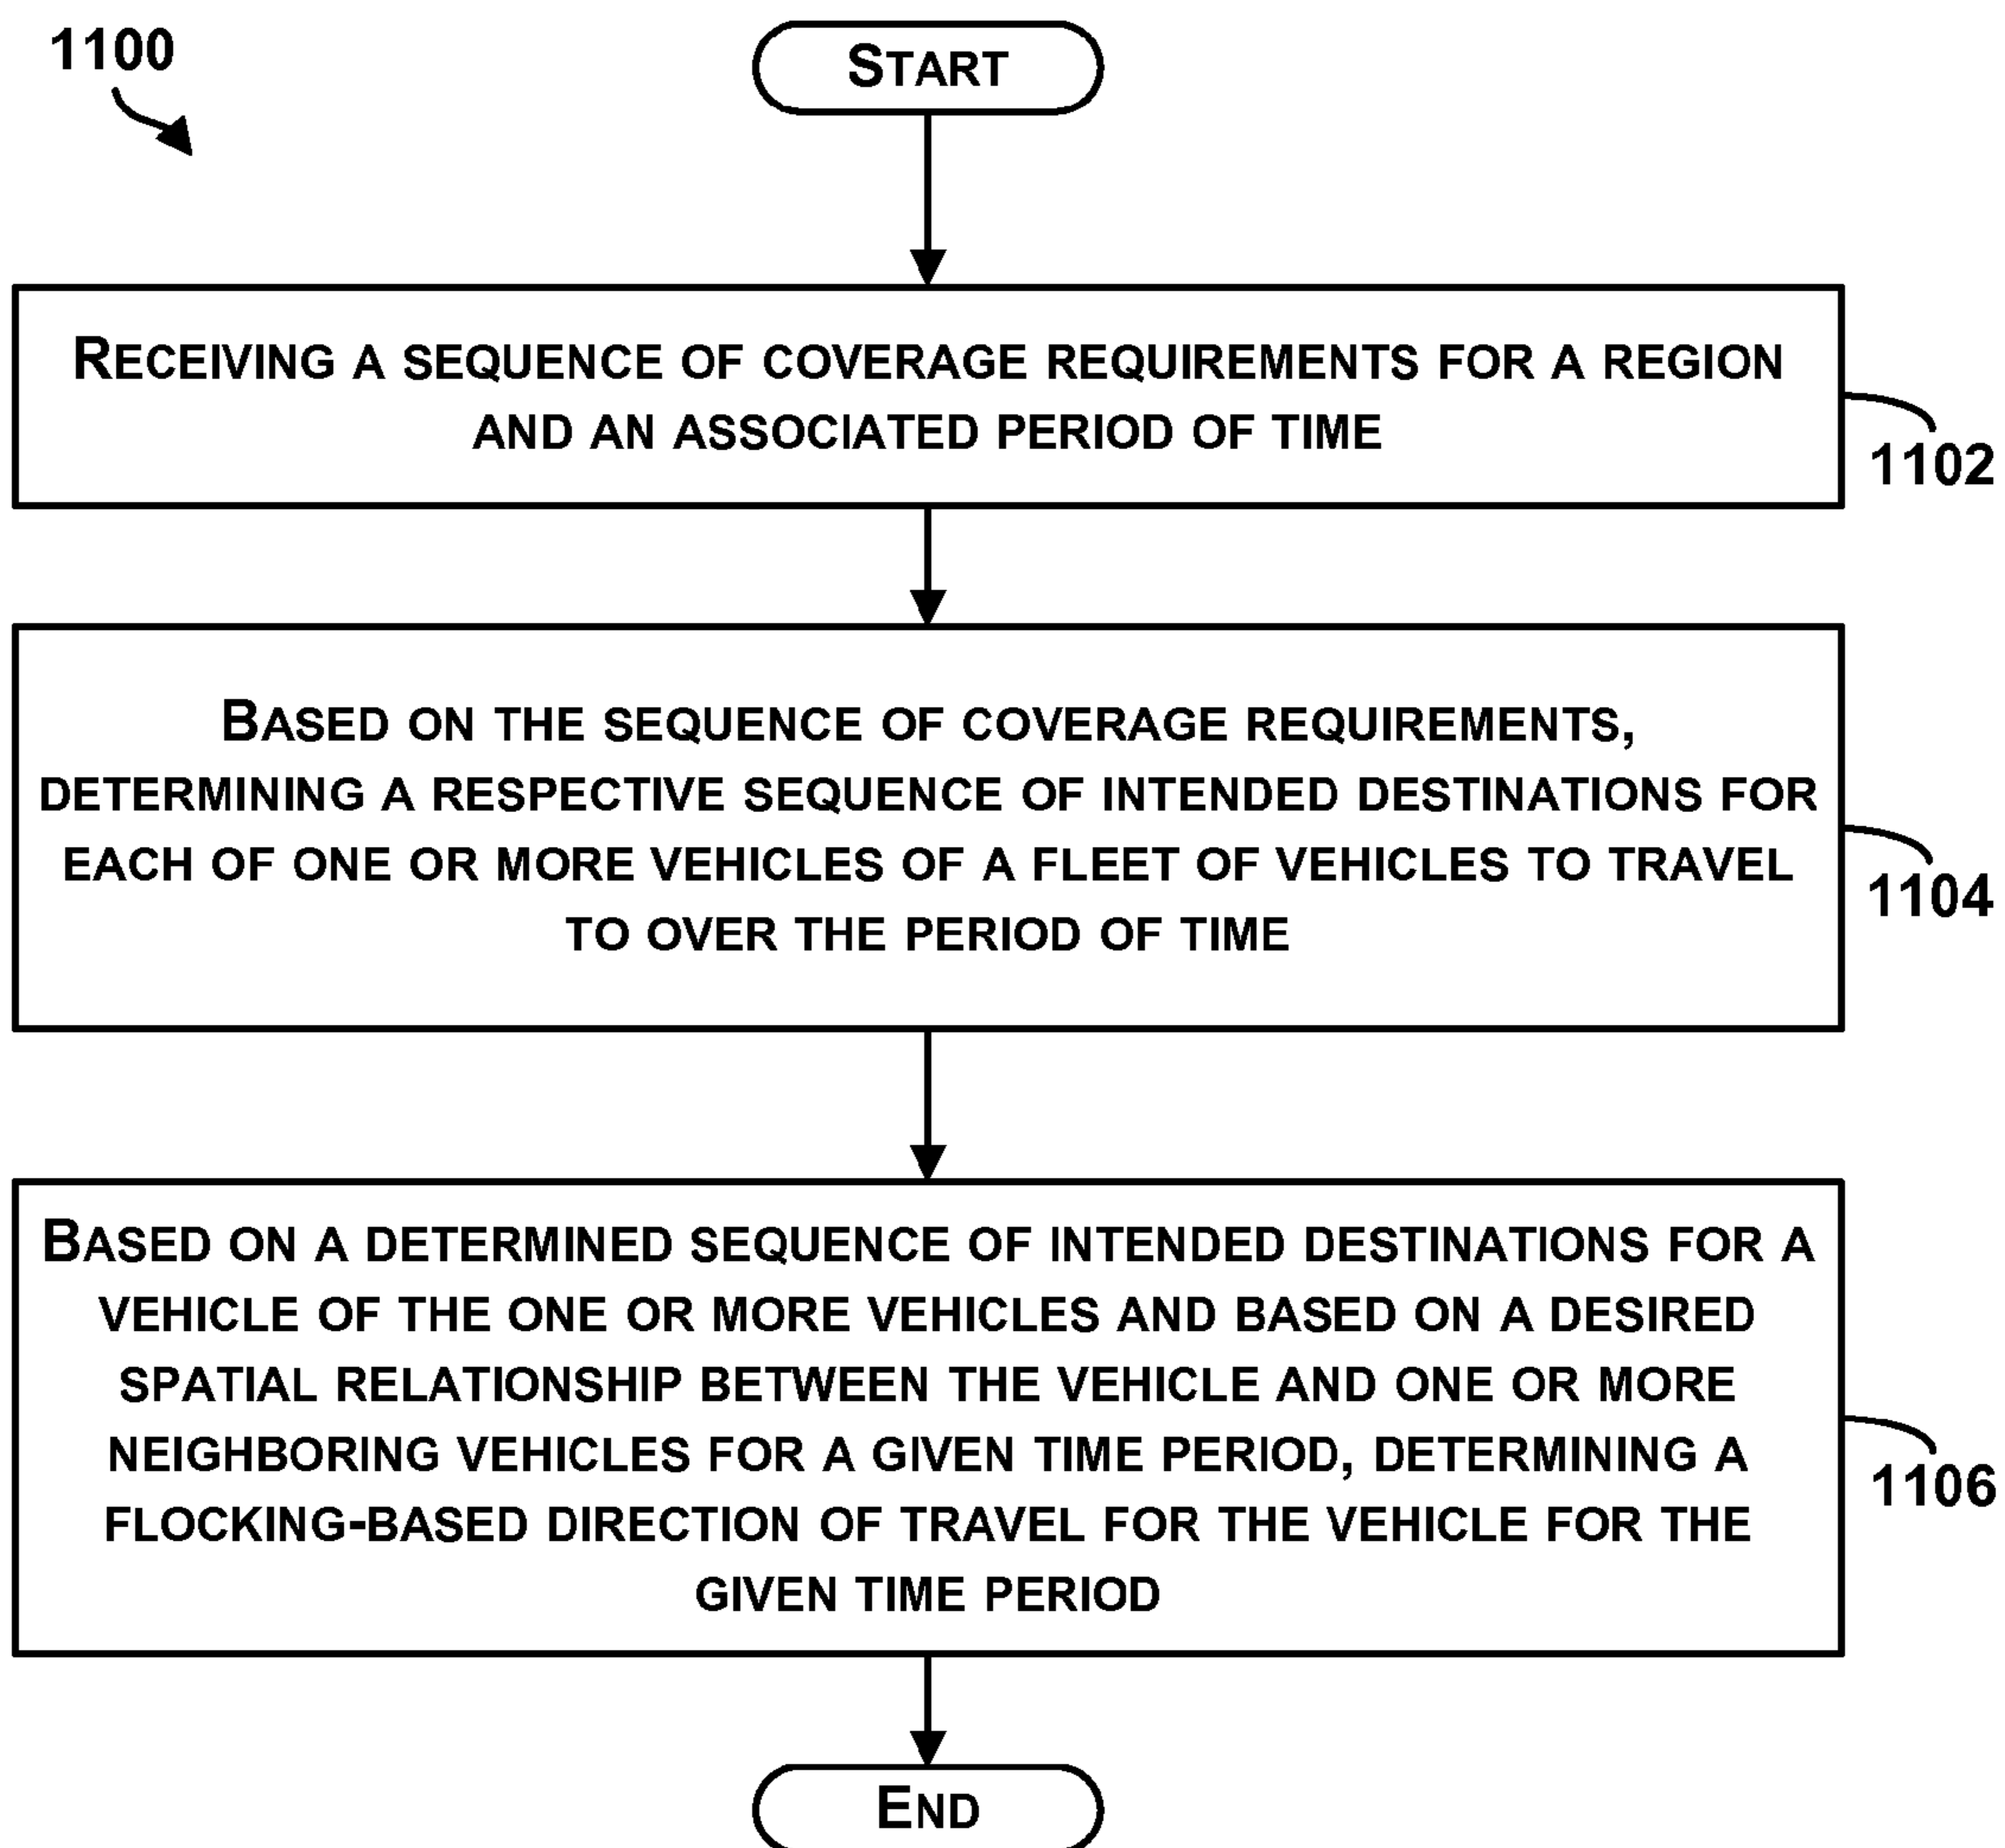
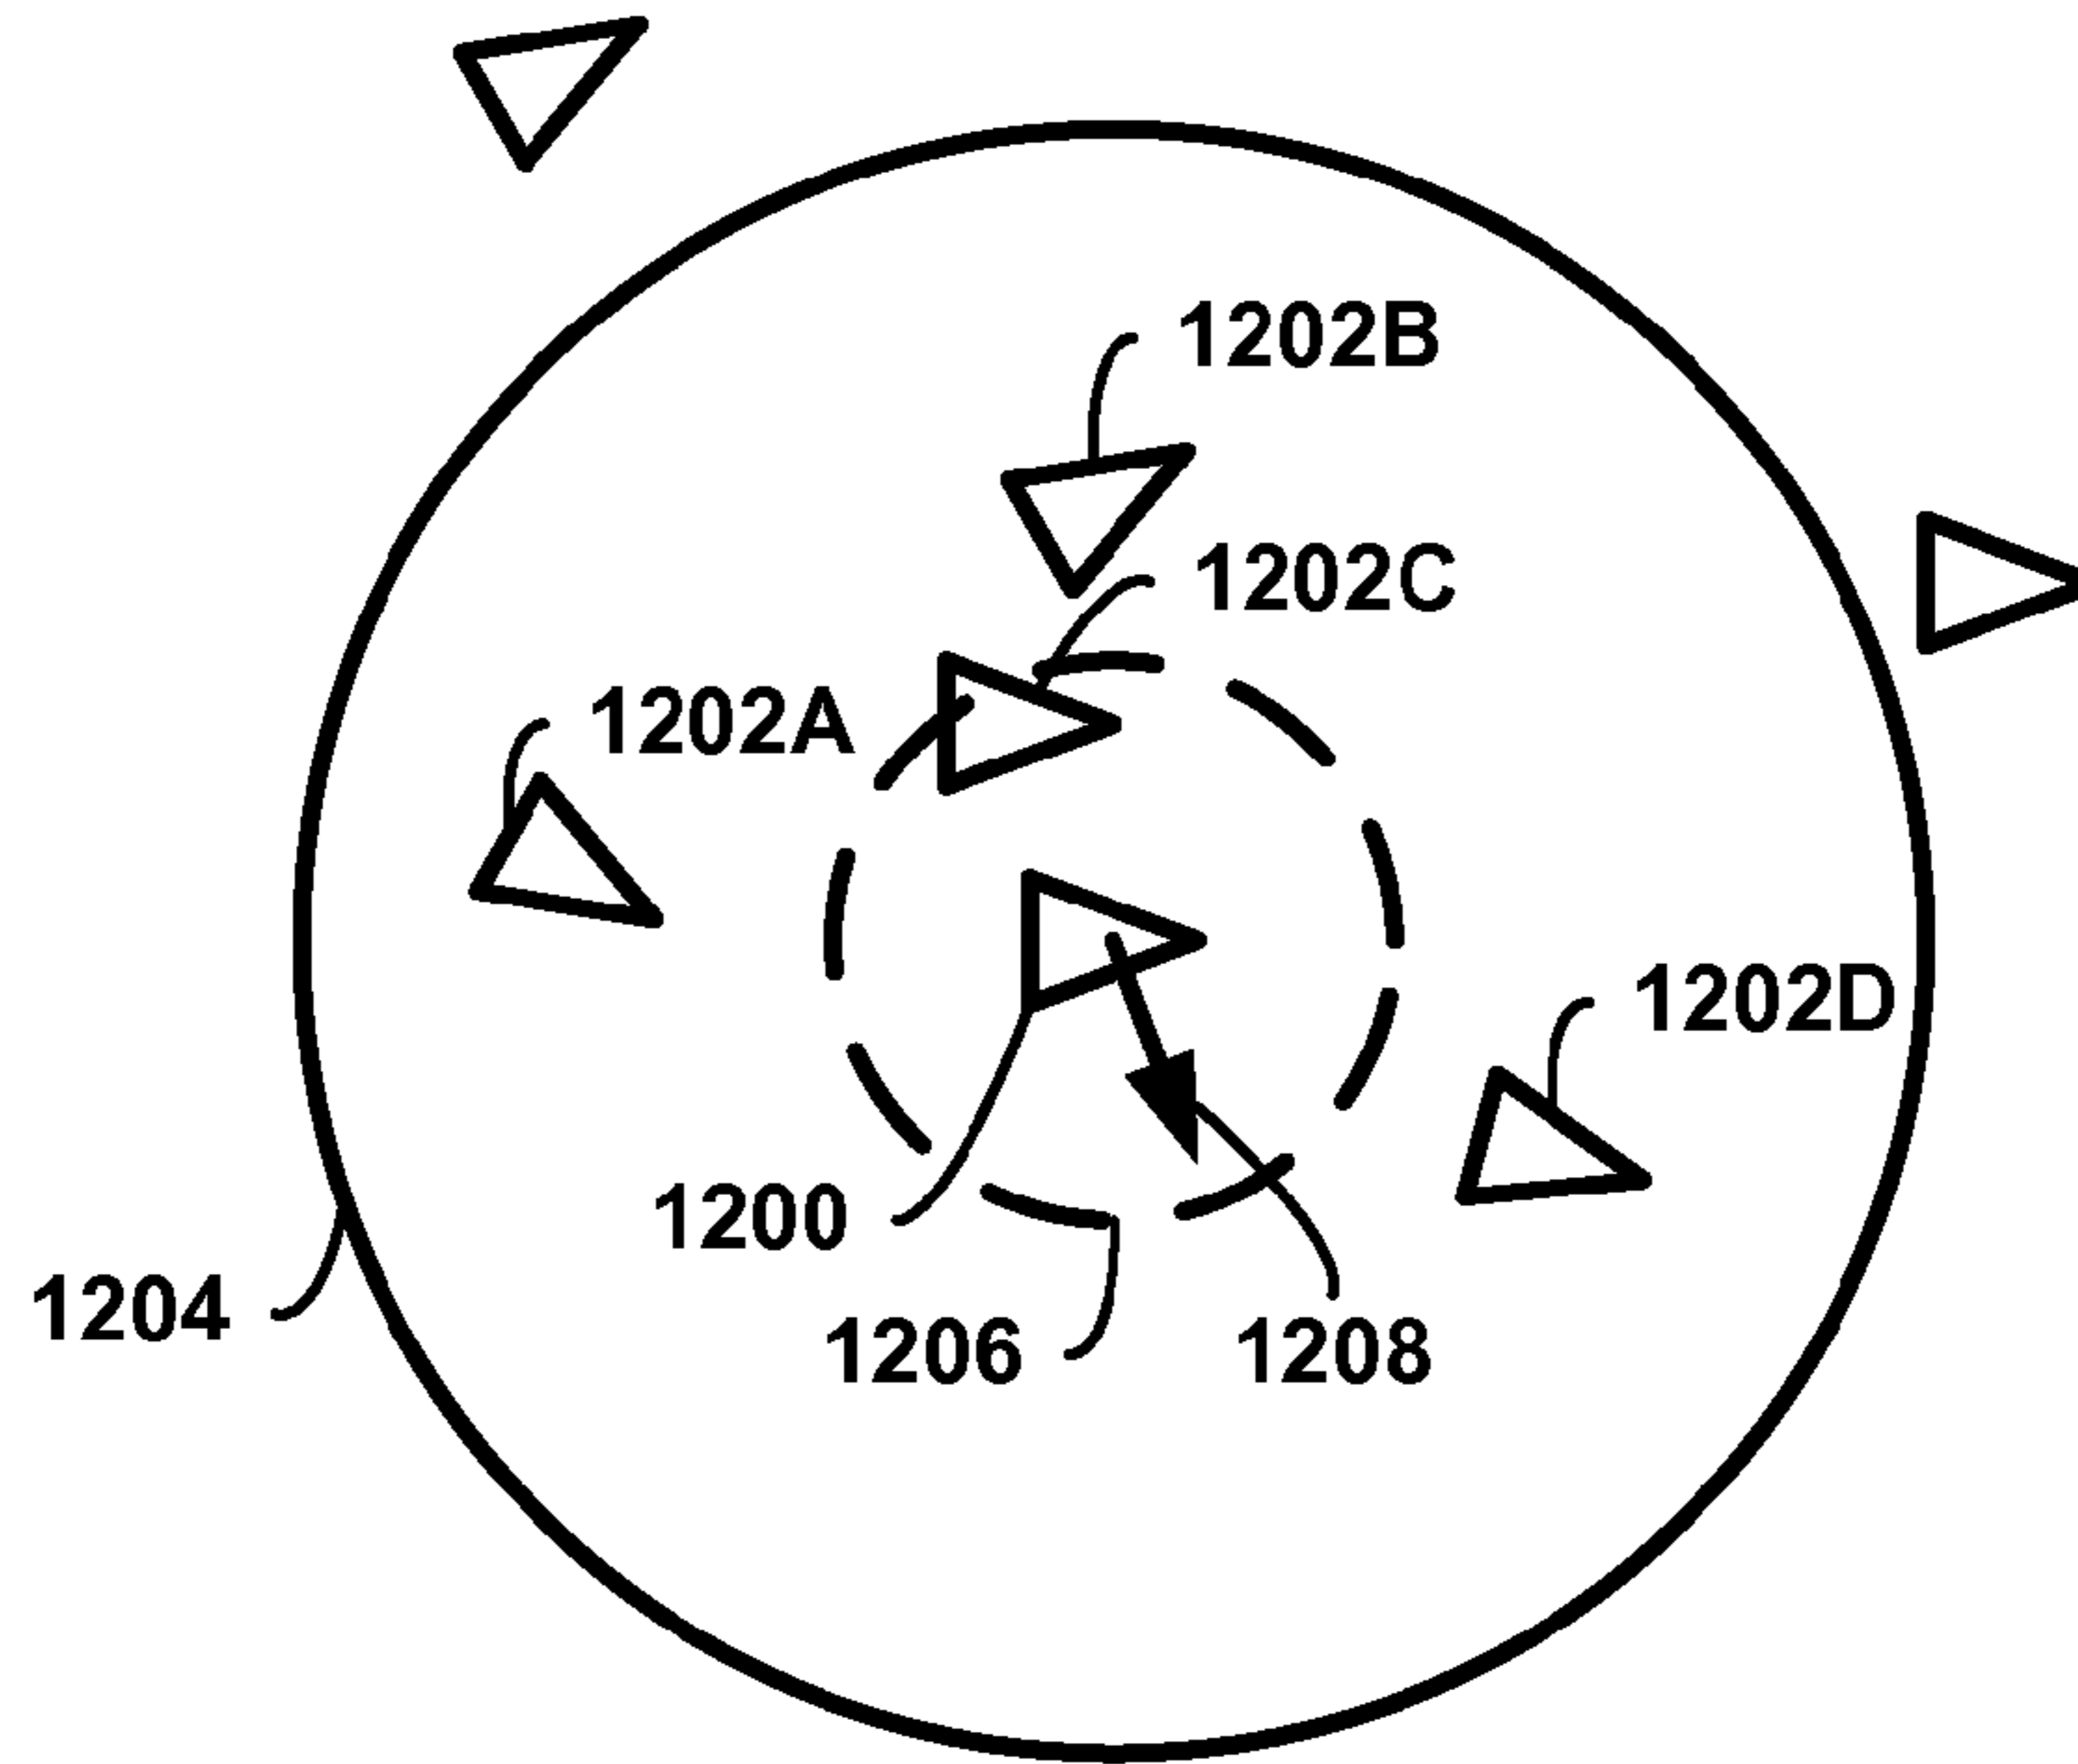
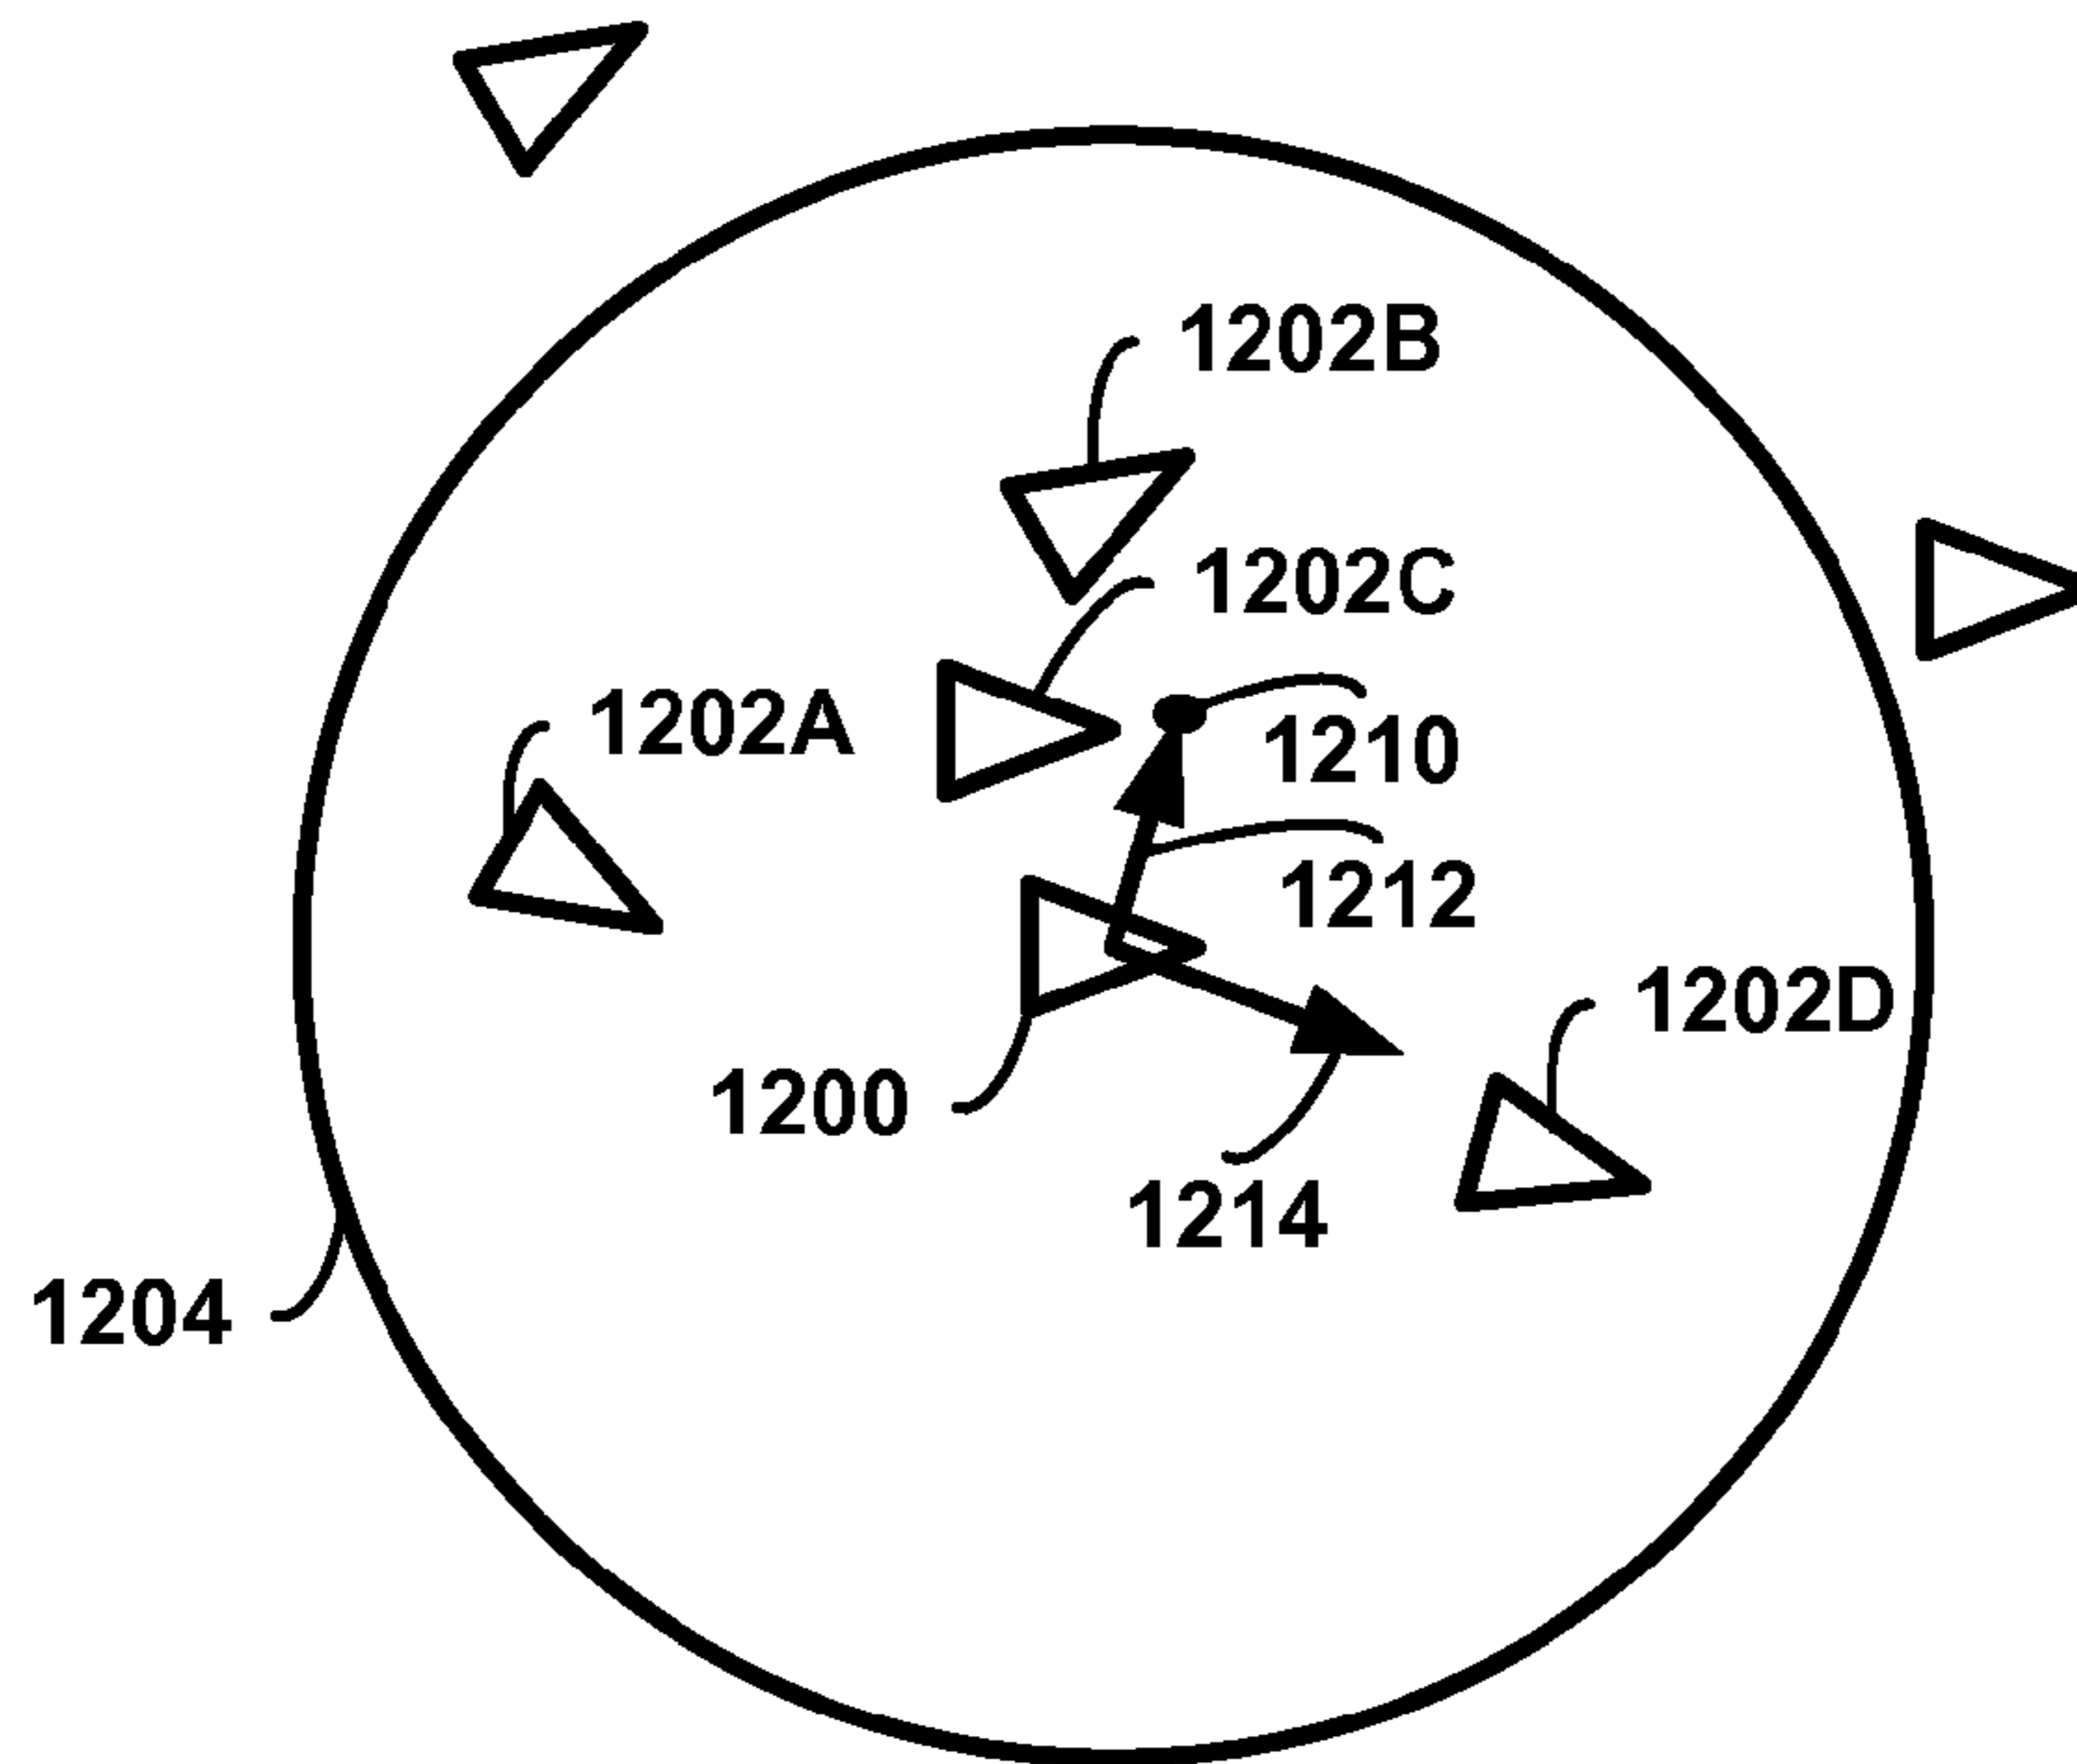


FIGURE 10

**FIGURE 11**

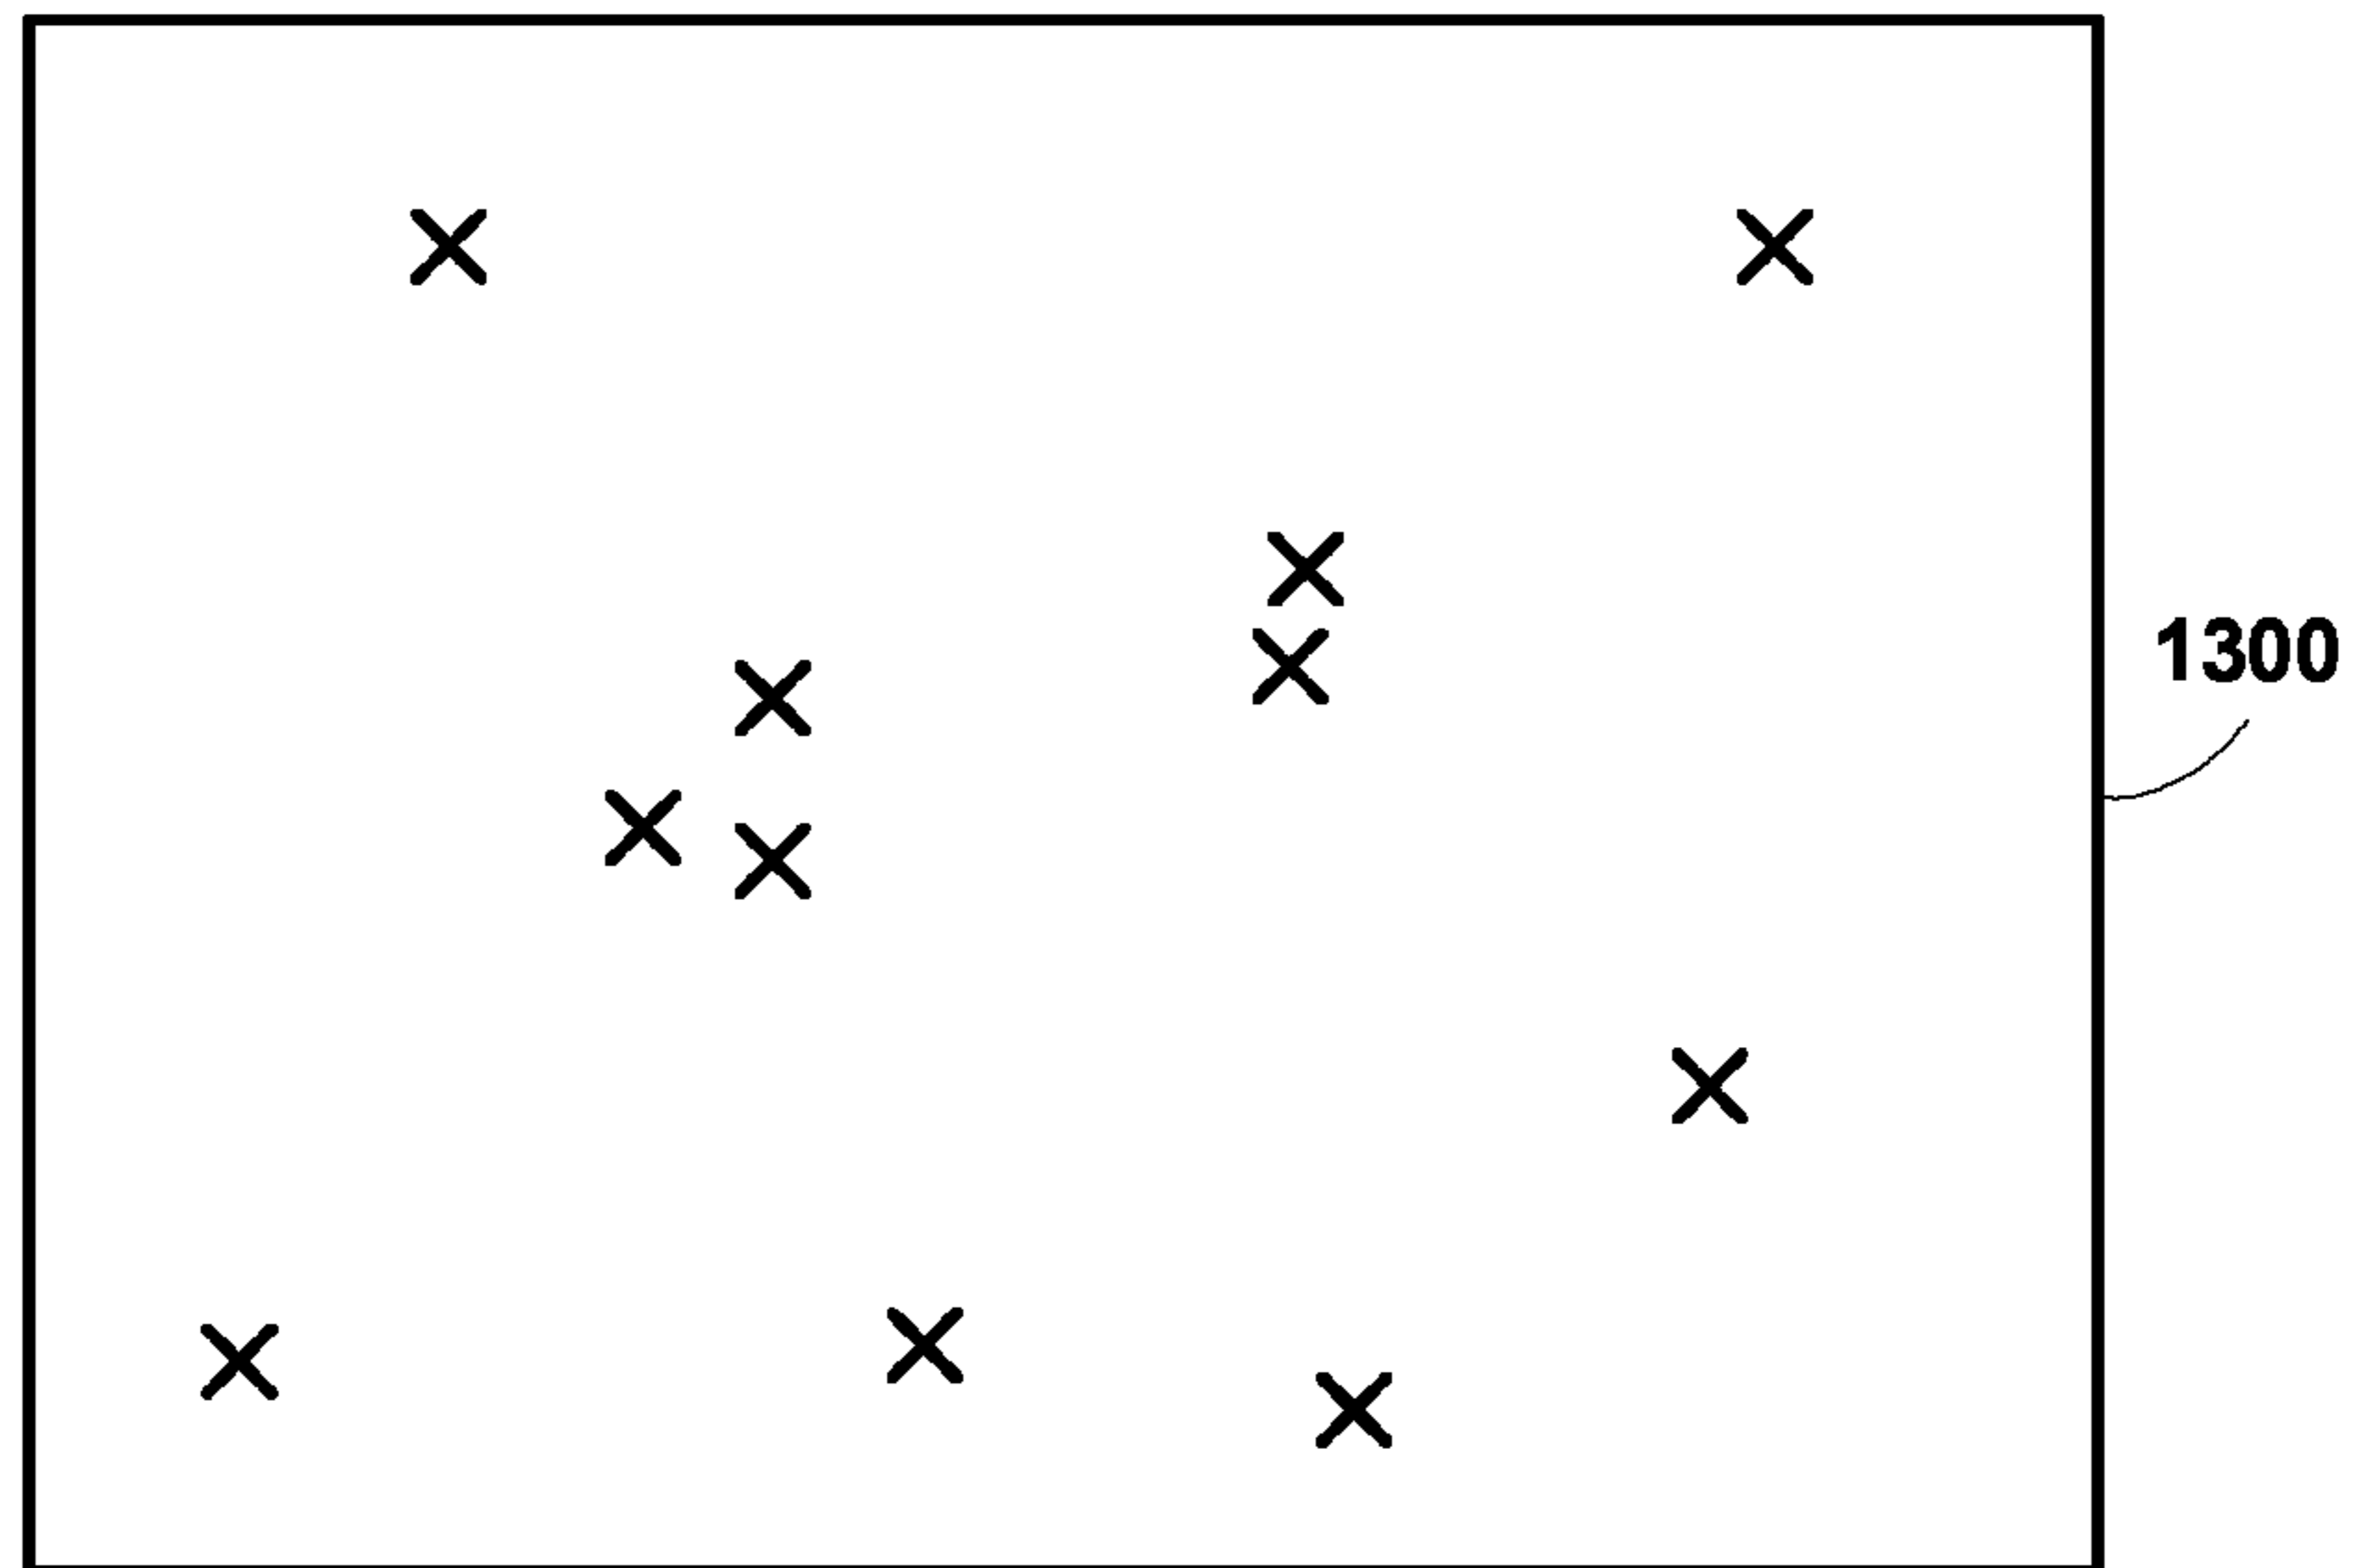


**FIGURE 12A**



**FIGURE 12B**

TIME: T = 2



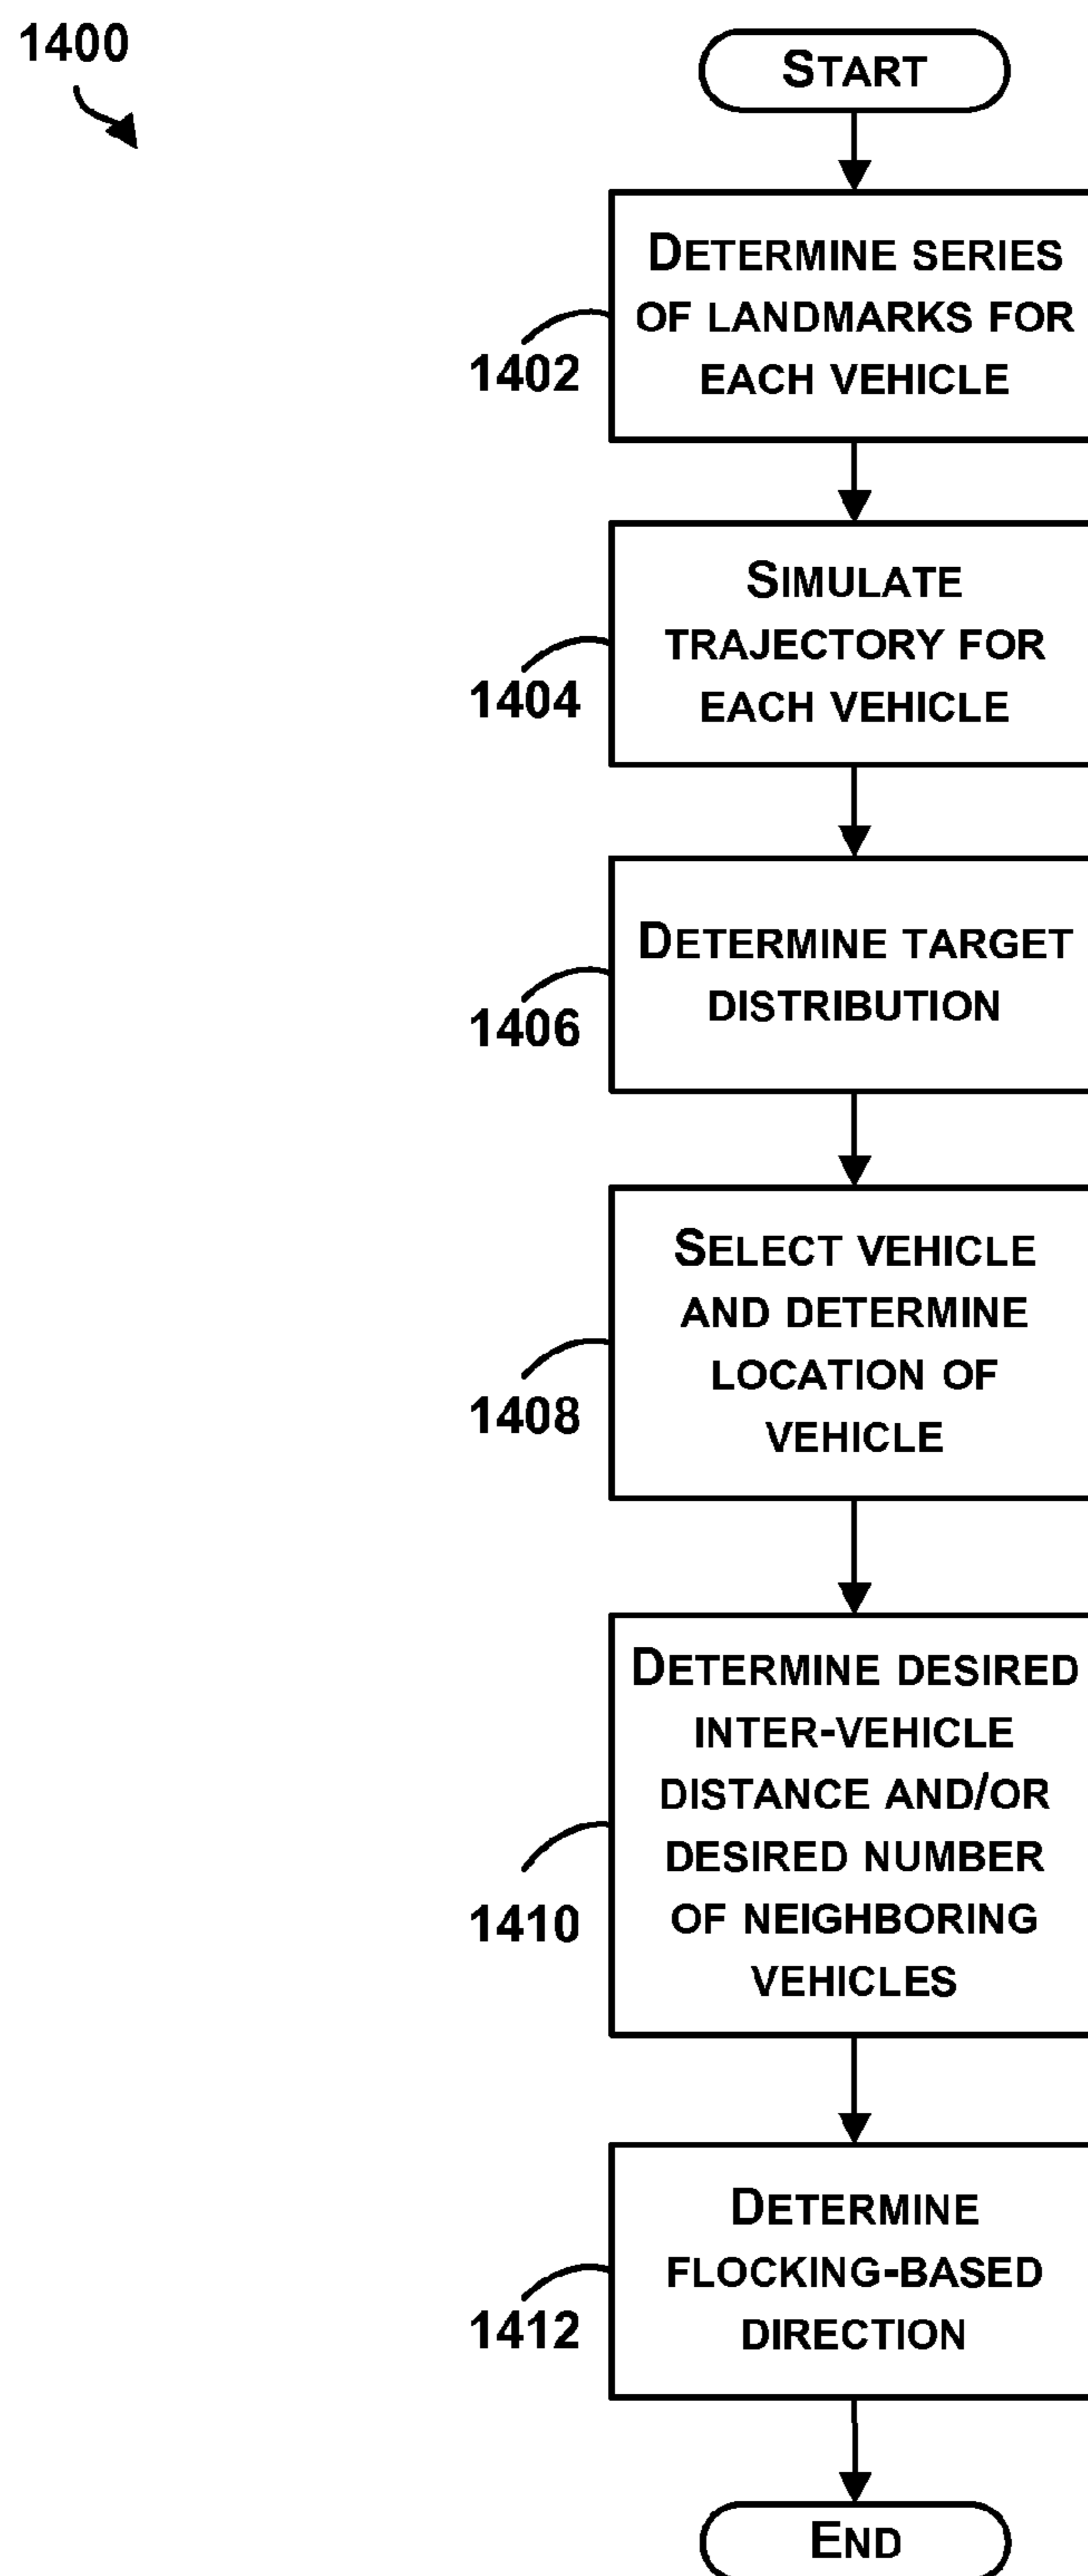
**FIGURE 13A**

TIME: T = 2

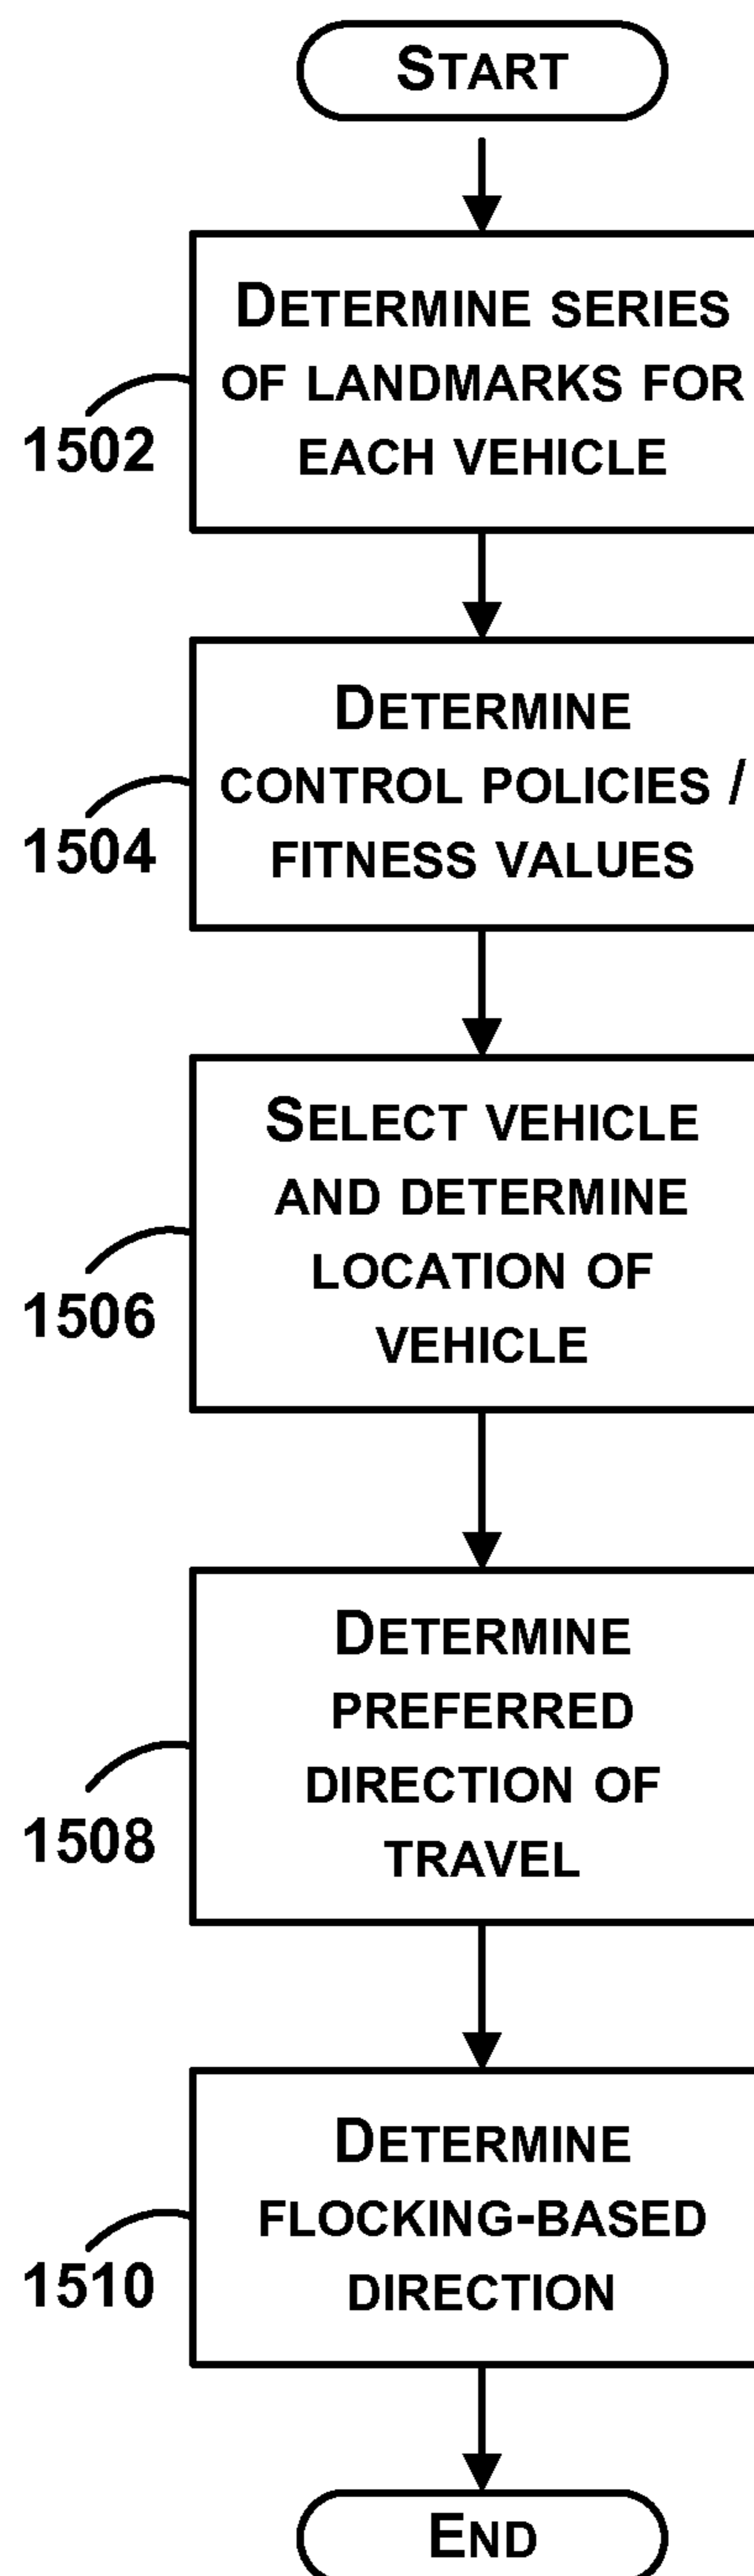
0.2	0.2	0.2	0	0	0.2	0.2	0.2
	X					X	
0	0.2	0.2	0	0.4	0	0.2	0
0	0.2	0.6	0.6	0.4	0.4	0	0
		X	X	X	X		
0	0.4	0.6	0.4	0.4	0	0.2	0
0.2	0	0.4	0.2	0	0.2	X	0.2
						0.2	
0.2	X	0.2	0.2	0.4	X	0.2	0.4
			X				0

**FIGURE 13B**



**FIGURE 14**

1500

**FIGURE 15**

1

## METHODS AND SYSTEMS FOR PERFORMING FLOCKING WHILE EXECUTING A LONG-RANGE FLEET PLAN

### BACKGROUND

Unless otherwise indicated herein, the materials described in this section are not prior art to the claims in this application and are not admitted to be prior art by inclusion in this section.

Computing devices such as personal computers, laptop computers, tablet computers, cellular phones, and countless types of Internet-capable devices are increasingly prevalent in numerous aspects of modern life. As such, the demand for data connectivity via the Internet, cellular data networks, and other such networks, is growing. However, there are many areas of the world where data connectivity is still unavailable, or if available, is unreliable and/or costly. Accordingly, additional network infrastructure is desirable.

### SUMMARY

In one example, a method is provided that comprises receiving a sequence of coverage requirements for a region and an associated period of time. An individual coverage requirement may indicate a desired distribution of vehicles among the region at a respective time period within the period of time. The method further comprises, based on the sequence of coverage requirements, determining, by a processor, a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time. Additionally, according to the method, based on a determined sequence of intended destinations for a vehicle of the one or more vehicles and based on a desired spatial relationship between the vehicle and one or more neighboring vehicles for a given time period, the processor may determine a flocking-based direction of travel for the vehicle for the given time period.

In another example, a computer-readable storage medium having stored therein instructions, that when executed by a computing device, cause the computing device to perform functions is provided. The functions comprise receiving a sequence of coverage requirements for a region and an associated period of time. An individual coverage requirement may indicate a desired distribution of vehicles among the region at a respective time period within the period of time. The functions further comprise, based on the sequence of coverage requirements, determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time. Additionally the functions comprise, based on a determined sequence of intended destinations for a vehicle of the one or more vehicles and based on a desired spatial relationship between the vehicle and one or more neighboring vehicles for a given time period, determining a flocking-based direction of travel for the vehicle for the given time period.

In still another example, a system is provided that comprises at least one processor, and data storage comprising program instructions executable by the at least one processor to cause the system to perform functions. The functions comprise receiving a sequence of coverage requirements for a region and an associated period of time. An individual coverage requirement may indicate a desired distribution of vehicles among the region at a respective phase within the period of time. The functions further comprise, based on the sequence of coverage requirements, determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of

2

time. Additionally the functions comprise, based on a determined sequence of intended destinations for a vehicle of the one or more vehicles and based on a desired spatial relationship between the vehicle and one or more neighboring vehicles for a given time period, determining a flocking-based direction of travel for the vehicle for the given time period.

These as well as other aspects, advantages, and alternatives, will become apparent to those of ordinary skill in the art by reading the following detailed description, with reference where appropriate to the accompanying figures.

### BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 is a block diagram illustrating an example balloon network.

FIG. 2 is a block diagram illustrating an example balloon-network control system.

FIG. 3 shows a high-altitude balloon according to an example embodiment.

FIG. 4 illustrates example trajectories for a balloon that is traveling from one of a plurality of starting landmarks to a landmark E.

FIG. 5 illustrates an example cost value map where location A is a starting location of a balloon, location B is a desired final location, and example cost values are assigned to cells surrounding location B.

FIG. 6 illustrates an area divided into cells, and for a given cell labeled X, example wind vectors  $A_1$ - $A_5$  that are determined based on wind conditions at five different altitude levels.

FIGS. 7A and 7B illustrate examples of determining a set of starting landmarks from which a balloon could reach a landmark E.

FIG. 8 illustrates an example graph of possible routes between a plurality of landmarks in accordance with an example embodiment.

FIGS. 9A and 9B conceptually illustrate an example determination of individual balloon trajectories.

FIG. 10 illustrates example trajectories for each of a plurality of balloons.

FIG. 11 is an example block diagram of a method to determine a flocking-based direction of travel, in accordance with at least some embodiments described herein.

FIGS. 12A and 12B are conceptual illustrations of flocking behaviors.

FIG. 13A illustrates estimated locations of each vehicle of a fleet of vehicles.

FIG. 13B illustrates a derived target distribution that is determined based on the estimated locations of FIG. 13A.

FIG. 14 is another example block diagram of a method to determine a flocking-based direction of travel, in accordance with at least some embodiments described herein.

FIG. 15 is still another example block diagram of a method to determine a flocking-based direction of travel, in accordance with at least some embodiments described herein.

### DETAILED DESCRIPTION

The following detailed description describes various features and functions of the disclosed systems and methods with reference to the accompanying figures. In the figures, similar symbols identify similar components, unless context dictates otherwise. The illustrative system and method embodiments described herein are not meant to be limiting. It may be readily understood that certain aspects of the dis-



closed systems and methods can be arranged and combined in a wide variety of different configurations, all of which are contemplated herein.

Examples herein may be configured to provide control of vehicles within a fleet of vehicles. The vehicles may be any number of types of vehicles including, for example, autonomous or semi-autonomous vehicles, such as aerial vehicles, boats, underwater vehicles, satellites, aerostats, etc. for mapping/surveying, sensing, providing connectivity, etc. In other examples, methods described herein may be used to control or dispatch autonomous cars to maintain a well-distributed fleet of autonomous taxis in a city to support rapid dispatch of taxis. The vehicles may be any type of mobile object, element, platform, etc.

In still further examples, vehicles may be or include elements of a data network that includes a plurality of balloons, for example, configured as a mesh network formed by high-altitude balloons deployed in the stratosphere. Since winds in the stratosphere may affect locations of the balloons in a differential manner, each balloon in an example network may be configured to change its horizontal position by adjusting its vertical position (i.e., altitude). For example, by adjusting its altitude, a balloon may be able find winds that will carry the balloon horizontally (e.g., latitudinally and/or longitudinally) to a desired horizontal location. In other examples, each balloon may be configured to change its horizontal position using propellers or other mechanisms to generate horizontal thrust rather than or in addition to adjusting its horizontal position by adjusting its vertical position.

Within examples, methods and systems are described with respect to elements of a fleet of vehicles being balloons in a fleet of balloons arranged as a data network. However, as mentioned above, methods and systems herein may be applicable to other types of vehicles of a fleet of vehicles, and thus, the examples described are not limiting.

When managing a fleet of autonomous balloons, a fleet administrator may often wish to satisfy multiple mission objectives, including, for example, both long-term goals and immediate goals. As an example of long-term goals, a respective sequence of intended destinations for one or more balloons may be determined (or received) based on a sequence of coverage requirements. The sequence of coverage requirements may specify desired amounts of balloon coverage for different parts of the world. For instance, the desired amounts of balloon coverage may include providing at least ten balloons in a first area, at least five balloons in a second area, and at least eighteen balloons in a third area at a first instance in time and then, at a second instance in time, providing at least five balloons in the first area, ten balloons in the second area, and twenty balloons in the third area. Thus, the sequence of coverage requirements may indicate desired geographic distributions of balloons at respective time periods.

In one example, satisfying long-term goals may involve determining a fleet plan that stages balloons in particular places ahead of deadlines such that the balloons will be able to arrive at a desired location on time may be required. Thus, planning a fleet plan to satisfy a sequence of coverage requirements may result in a sequence of way points that each balloon should pass through.

It may also be desirable to evaluate or plan for the distribution of the balloons as each of the balloons pass between waypoints. As an example of immediate goals, an administrator might desire the fleet of balloons to not bunch up, to not have any point in space be more than a set distance from at least N balloons (e.g., in order to provide good sensor coverage, or to provide continuous wireless communication coverage), or to have every balloon be within a set distance of at

least N other balloons (e.g., in order to maintain a balloon-to-balloon communications network). These types of immediate goals are often amenable to reactive, real-time control systems, such as flocking algorithms.

Described herein are example mechanisms (e.g., systems and methods) for managing both long-term goals and immediate goals when coordinating the trajectories for a fleet of autonomous or semi-autonomous vehicles.

In one example, a system may use a long-term planner to derive a plan for a fleet of vehicles that satisfies long-term goals. The system may then simulate the fleet executing the long-term plan, resulting in a detailed trajectory for each vehicle. The system may also be configured to determine, for any time within the long-term plan, a derived target distribution. The derived target distribution may be determined based on the location of each vehicle at that time in the detailed trajectory for the vehicle. The system may then use a flocking algorithm that incorporates information associated with a target distribution to determine a flocking-based direction of travel for each vehicle.

The term flocking as used herein may generally refer to any type of behavior that is determined based on immediate circumstances, as opposed to global or long-term goals. Thus, a flocking algorithm may be any algorithm that locally solves for a direction and/or behavior based on one or more immediate circumstances of a vehicle or group of vehicles. The immediate circumstances may include, for example, the location and/or behavior of a vehicle and one or more neighboring vehicles as further described below. Other examples may also exist, and the example is not meant to be limiting in any way.

In another example, a system may use a long-term planner to derive a plan for a fleet of vehicles that satisfies long-term goals. From the long-term plan, the system may derive a control policy to enable each vehicle to satisfy a series of goals. For example, the control policy may specify for each vehicle and each location, the preferred direction of travel in order to get as close as possible to an intended goal. The system may then use a flocking algorithm that allows a preferred direction to be incorporated into each vehicle's flocking-based control choice to determine a flocking-based direction of travel for each vehicle. Other example systems and methods are also contemplated and described hereinafter.

Referring now to the Figures, FIG. 1 is a block diagram illustrating an example balloon network 100. As shown, the balloon network 100 includes balloons 102A to 102F, which are configured to communicate with one another via free-space optical links 104 and/or RF links 108. Configured as such, balloons 102A to 102F may collectively function as a mesh network for packet-data communications. Further, at least some of balloons 102A and 102B may be configured for RF communications with ground-based stations 106 via respective RF links 108. Yet further, some balloons, such as balloon 102F, may be configured to communicate via optical link 110 with ground-based station 112.

In an exemplary embodiment, balloons 102A to 102F are high-altitude balloons, which are deployed in the stratosphere. At moderate latitudes, the stratosphere includes altitudes between approximately 10 kilometers (km) and 50 km altitude above the surface. At the poles, the stratosphere starts at an altitude of approximately 8 km. In an exemplary embodiment, high-altitude balloons may be generally configured to operate in an altitude range within the stratosphere that has a wide range of wind speeds. For instance, between 17 km and 20 km altitude above the surface, the average wind speed may be between about 30 mph to about 35 mph, while the max wind speed may be up to 200 mph. As another example, at about 26 km altitude above the surface, the aver-



age wind speed may be between about 60 mph to about 65 mph, while the max wind speed may exceed 300 mph.

More specifically, in a high-altitude-balloon network, balloons **102A** to **102F** may generally be configured to operate at altitudes between 18 km and 25 km (although other altitudes are possible). At such altitudes, the winds may vary with latitude and by season, and the variations can be modeled in a reasonably accurate manner. Additionally, altitudes above 18 km are typically above the maximum flight level designated for commercial air traffic. Therefore, interference with commercial flights is not a concern when balloons are deployed between 18 km and 25 km.

To transmit data to another balloon, a given balloon **102A** to **102F** may be configured to transmit an optical signal via an optical link **104**. In an exemplary embodiment, a given balloon **102A** to **102F** may use one or more high-power light-emitting diodes (LEDs) to transmit an optical signal. Alternatively, some or all of balloons **102A** to **102F** may include laser systems for free-space optical communications over optical links **104**. Other types of free-space optical communication are possible. Further, in order to receive an optical signal from another balloon via an optical link **104**, a given balloon **102A** to **102F** may include one or more optical receivers.

In a further aspect, balloons **102A** to **102F** may utilize one or more of various different RF air-interface protocols for communication with ground-based stations **106** via respective RF links **108**. For instance, some or all of balloons **102A** to **102F** may be configured to communicate with ground-based station **106** using protocols described in IEEE 802.11 (including any of the IEEE 802.11 revisions), various cellular protocols such as GSM, CDMA, UMTS, EV-DO, WiMAX, and/or LTE, and/or one or more propriety protocols developed for balloon-ground RF communication, among other possibilities.

In a further aspect in balloon network **100**, balloon **102F** is configured as a downlink balloon to provide a high-capacity air-ground link. Like other balloons in an exemplary network, a downlink balloon **102F** may be operable for optical communication with other balloons via optical links **104**. However, a downlink balloon **102F** may also be configured for free-space optical communication with a ground-based station **112** via an optical link **110**. Optical link **110** may therefore serve as a high-capacity link (as compared to an RF link **108**) between the balloon network **100** and ground-based station **112**.

Note that in some implementations, a downlink balloon **102F** may additionally be operable for RF communication with ground-based station **106**. In other cases, a downlink balloon **102F** may only use an optical link for balloon-to-ground communications. Further, while the arrangement shown in FIG. 1 includes just one downlink balloon **102F**, an exemplary balloon network can also include multiple downlink balloons. On the other hand, a balloon network can also be implemented without any downlink balloons.

In other implementations, a downlink balloon may be equipped with a specialized, high-bandwidth RF communication system for balloon-to-ground communications, instead of, or in addition to, a free-space optical communication system. The high-bandwidth RF communication system may take the form of an ultra-wideband system, which may provide an RF link with substantially the same capacity as one of the optical links **104**.

Ground-based stations, such as ground-based station **106** and/or **112**, may include components such as transceivers, transmitters, and/or receivers for communication via RF links and/or optical links with a balloon network. Further, a

ground-based station may use various air-interface protocols in order to communicate with a balloon **102A** to **102F** over an RF link. As such, a ground-based station **106** may be configured as an access point via which various devices can connect to balloon network **100**.

In a further aspect, some or all balloons **102A** to **102F** could be configured to establish a communication link with space-based satellites in addition to, or as an alternative to, a ground-based communication link. In some embodiments, a balloon may communicate with a satellite via an optical link.

Further, some ground-based stations, such as ground-based station **106**, may be configured as gateways between balloon network **100** and one or more other networks. Such a ground-based station **106** may thus serve as an interface between the balloon network and the Internet, a cellular service provider's network, and/or other types of networks.

As noted, balloons **102A** to **102F** may collectively function as a mesh network. More specifically, since balloons **102A** to **102F** may communicate with one another using free-space optical links, the balloons may collectively function as a free-space optical mesh network.

In a mesh-network configuration, each balloon **102A** to **102F** may function as a node of the mesh network, which is operable to receive data directed to the balloon and to route data to other balloons. As such, data may be routed from a source balloon to a destination balloon by determining an appropriate sequence of optical links between the source balloon and the destination balloon. These optical links may be collectively referred to as a "lightpath" for the connection between the source and destination balloons. Further, each of the optical links may be referred to as a "hop" on the lightpath.

To operate as a mesh network, balloons **102A** to **102F** may employ various routing techniques and self-healing algorithms. In some embodiments, a balloon network **100** may employ adaptive or dynamic routing, where a lightpath between a source and destination balloon is determined and set-up when the connection is needed, and released at a later time. Further, when adaptive routing is used, the lightpath may be determined dynamically depending upon the current state, past state, and/or predicted state of the balloon network.

In addition, the network topology may change as the balloons **102A** to **102F** move relative to one another and/or relative to the ground. Accordingly, an exemplary balloon network **100** may apply a mesh protocol to update the state of the network as the topology of the network changes. For example, to address the mobility of the balloons **102A** to **102F**, balloon network **100** may employ and/or adapt various techniques that are employed in mobile ad hoc networks (MANETs). Other examples are possible as well.

In some implementations, a balloon network **100** may be configured as a transparent mesh network. For instance, the balloons may include components for physical switching that is entirely optical, without any electrical components involved in the physical routing of optical signals. Thus, in a transparent configuration with optical switching, signals travel through a multi-hop lightpath that is entirely optical.

In other implementations, the balloon network **100** may implement a free-space optical mesh network that is opaque. In an opaque configuration, some or all balloons **102A** to **102F** may implement optical-electrical-optical (OEO) switching. For example, some or all balloons may include optical cross-connects (OXC) for OEO conversion of optical signals. Other opaque configurations are also possible.

In a further aspect, balloons in an exemplary balloon network **100** may implement wavelength division multiplexing (WDM), which may help to increase link capacity. When WDM is implemented with transparent switching, physical



lightpaths through the balloon network may be subject to the “wavelength continuity constraint.” More specifically, because the switching in a transparent network is entirely optical, the same wavelength can be assigned for all optical links on a given lightpath.

An opaque configuration, on the other hand, may avoid the wavelength continuity constraint. In particular, balloons in an opaque balloon network may include the OEO switching systems operable for wavelength conversion. As a result, balloons can convert the wavelength of an optical signal at each hop along a lightpath.

Further, various routing algorithms may be employed in an opaque configuration. For example, to determine a primary lightpath and/or one or more diverse backup lightpaths for a given connection, exemplary balloons may apply or consider shortest-path routing techniques such as Dijkstra’s algorithm and k-shortest path, and/or edge and node-diverse or disjoint routing such as Suurballe’s algorithm, among others. Additionally or alternatively, techniques for maintaining a particular quality of service (QoS) may be employed when determining a lightpath. Other techniques are also possible.

In some examples, a balloon network **100** may have a non-uniform topology. For instance, exemplary embodiments may involve topologies where balloons are distributed more or less densely in certain areas, for various reasons. As an example, to help meet the higher bandwidth demands that are typical in urban areas, balloons may be clustered more densely over urban areas. For similar reasons, the distribution of balloons may be more dense over land than over large bodies of water. Many other examples of non-uniform topologies are possible.

FIG. 2 is a block diagram illustrating an example balloon-network control system. In particular, FIG. 2 shows a distributed control system, which includes a central control system **200** and a number of regional control-systems **202A** to **202C**. Such a control system may be configured to coordinate certain functionality for balloon network **204**, and as such, may be configured to control and/or coordinate certain functions for balloons **206A** to **206I**.

In the illustrated embodiment, central control system **200** may be configured to communicate with balloons **206A** to **206I** via a number of regional control systems **202A** to **202C**. These regional control systems **202A** to **202C** may be configured to receive communications and/or aggregate data from balloons in the respective geographic areas that the balloons cover, and to relay the communications and/or data to central control system **200**. Further, regional control systems **202A** to **202C** may be configured to route communications from central control system **200** to the balloons in their respective geographic areas. For instance, as shown in FIG. 2, regional control system **202A** may relay communications and/or data between balloons **206A** to **206C** and central control system **200**, regional control system **202B** may relay communications and/or data between balloons **206D** to **206F** and central control system **200**, and regional control system **202C** may relay communications and/or data between balloons **206G** to **206I** and central control system **200**.

In order to facilitate communications between the central control system **200** and balloons **206A** to **206I**, certain balloons may be configured as downlink balloons, which are operable to communicate with regional control systems **202A** to **202C**. Accordingly, each regional control system **202A** to **202C** may be configured to communicate with the downlink balloon or balloons in the respective geographic area it covers. For example, in the illustrated embodiment, balloons **204A**, **204D**, and **204H** are configured as downlink balloons. As such, regional control systems **202A** to **202C** may respec-

tively communicate with balloons **204A**, **204D**, and **204H** via optical links **206**, **208**, and **210**, respectively.

In a centralized control arrangement, such as that shown in FIG. 2, the central control system **200** (and possibly regional control systems **202A** to **202C** as well) may coordinate certain mesh-networking functions for balloon network **204**. For example, balloons **206A** to **206I** may send the central control system **200** certain state information, which the central control system **200** may utilize to determine the state of balloon network **204**. The state information from a given balloon may include location data, optical-link information (e.g., the identity of other balloons with which the balloon has established an optical link, the bandwidth of the link, wavelength usage and/or availability on a link, etc.), wind data collected by the balloon, and/or other types of information. Accordingly, the central control system **200** may aggregate state information from some or all of the balloons **206A** to **206I** in order to determine an overall state of the network.

The overall state of the network may then be used to coordinate and/or facilitate certain mesh-networking functions such as determining lightpaths for connections. For example, the central control system **200** may determine a current topology based on the aggregate state information from some or all of the balloons **206A** to **206I**. The topology may provide a picture of the current optical links that are available in balloon network and/or the wavelength availability on the links. This topology may then be sent to some or all of the balloons so that a routing technique may be employed to select appropriate lightpaths (and possibly backup lightpaths) for communications through the balloon network **204**.

In some embodiments, control and coordination of a balloon network may be shared by a ground-based control system and a balloon network to varying degrees, depending upon the implementation. In fact, in some embodiments, there may be no ground-based control systems. In such an embodiment, all network control and coordination functions may be implemented by the balloon network itself. For example, certain balloons may be configured to provide the same or similar functions as central control system **200** and/or regional control systems **202A** to **202C**. Other examples are also possible.

Furthermore, control and/or coordination of a balloon network may be de-centralized. For example, each balloon may relay state information to, and receive state information from, some or all nearby balloons. Further, each balloon may relay state information that it receives from a nearby balloon to some or all nearby balloons. When all balloons do so, each balloon may be able to individually determine the state of the network. Alternatively, certain balloons may be designated to aggregate state information for a given portion of the network. These balloons may then coordinate with one another to determine the overall state of the network.

Various types of balloon systems may be incorporated in an exemplary balloon network. As noted above, an exemplary embodiment may utilize high-altitude balloons, which typically operate in an altitude range between 18 km and 25 km. FIG. 3 shows a high-altitude balloon **300**, according to an exemplary embodiment. As shown, the balloon **300** includes an envelope **302**, a skirt **304**, a payload **306**, and a cut-down system **308** that is attached between the balloon **302** and payload **306**.

The envelope **302** and skirt **304** may take various forms, for instance, the envelope **302** and/or skirt **304** may be made of metalized Mylar or BoPet. Alternatively or additionally, some or all of the envelope **302** and/or skirt **304** may be constructed from a highly-flexible latex material or a rubber material such as chloroprene. Further, the shape and size of the envelope



302 and skirt 304 may vary depending upon the particular implementation. Additionally, the envelope 302 may be filled with various different types of gases, such as helium and/or hydrogen.

The payload 306 of balloon 300 may include a processor 312 and on-board data storage, such as memory 314. The memory 314 may take the form of or include a non-transitory computer-readable medium. The non-transitory computer-readable medium may have instructions stored thereon, which can be accessed and executed by the processor 312 in order to carry out the balloon functions described herein.

The payload 306 of balloon 300 may also include various other types of equipment and systems to provide a number of different functions. For example, payload 306 may include optical communication system 316, which may transmit optical signals via an ultra-bright LED system 320, and which may receive optical signals via an optical-communication receiver 322 (e.g., a photo-diode receiver system). Further, payload 306 may include an RF communication system 318, which may transmit and/or receive RF communications via an antenna system 340. The payload may further include a positioning system 324 that can include a GPS or other location-based sensors.

The payload 306 may also include a power supply 326 to supply power to the various components of balloon 300. The power supply 326 may include or take the form of a rechargeable battery. In addition, the balloon 300 may include a solar power generation system 327. The solar power generation system 327 may include solar panels and could be used to generate power that charges and/or is distributed by the power supply 326.

Further, payload 306 may include various types of other systems 328 and sensors 334. For example, payload 306 may include one or more video and/or still cameras, a GPS system, various motion sensors (e.g., accelerometers, gyroscopes, and/or compasses), and/or various sensors for capturing environmental data. Further, some or all of the components within payload 306 may be implemented in a radiosonde, which may be operable to measure, e.g., pressure, altitude, geographical position (latitude and longitude), temperature, relative humidity, and/or wind speed and/or direction, among other information.

As noted, balloon 300 includes an ultra-bright LED system 320 for free-space optical communication with other balloons. As such, optical communication system 316 may be configured to transmit a free-space optical signal by modulating the ultra-bright LED system 320. The optical communication system 316 may be implemented with mechanical systems and/or with hardware, firmware, and/or software.

In a further aspect, balloon 300 may be configured for altitude control. For instance, balloon 300 may include a variable buoyancy system, which is configured to change the altitude of the balloon 300 by adjusting the volume and/or density of the gas in the balloon 300. A variable buoyancy system may take various forms, and may generally be any system that can change the volume and/or density of gas in envelope 302.

In an exemplary embodiment, a variable buoyancy system may include a bladder 310 that is located inside of envelope 302. The bladder 310 could be an elastic chamber configured to hold liquid and/or gas. Alternatively, the bladder 310 need not be inside the envelope 302. For instance, the bladder 310 could be a rigid bladder that could be pressurized well beyond neutral pressure. The buoyancy of the balloon 300 may therefore be adjusted by changing the density and/or volume of the gas in bladder 310. To change the density in bladder 310, balloon 300 may be configured with systems and/or mecha-

nisms for heating and/or cooling the gas in bladder 310. Further, to change the volume, balloon 300 may include pumps or other features for adding gas to and/or removing gas from bladder 310. Additionally or alternatively, to change the volume of bladder 310, balloon 300 may include release valves or other features that are controllable to allow gas to escape from bladder 310. Multiple bladders 310 could be implemented within the scope of this disclosure. For instance, multiple bladders could be used to improve balloon stability.

In an example embodiment, the envelope 302 could be filled with helium, hydrogen or other lighter-than-air material. The envelope 302 could thus have an associated upward buoyancy force. In such an embodiment, air in the bladder 310 could be considered a ballast tank that may have an associated downward ballast force. In another example embodiment, the amount of air in the bladder 310 could be changed by pumping air (e.g., with an air compressor) into and out of the bladder 310. By adjusting the amount of air in the bladder 310, the ballast force may be controlled. In some embodiments, the ballast force may be used, in part, to counteract the buoyancy force and/or to provide altitude stability.

In other embodiments, the envelope 302 could be substantially rigid and include an enclosed volume. Air could be evacuated from envelope 302 while the enclosed volume is substantially maintained. In other words, at least a partial vacuum could be created and maintained within the enclosed volume. Thus, the envelope 302 and the enclosed volume could become lighter than air and provide a buoyancy force. In yet other embodiments, air or another material could be controllably introduced into the partial vacuum of the enclosed volume in an effort to adjust the overall buoyancy force and/or to provide altitude control.

In another embodiment, a portion of the envelope 302 could be a first color (e.g., black) and/or a first material from the rest of envelope 302, which may have a second color (e.g., white) and/or a second material. For instance, the first color and/or first material could be configured to absorb a relatively larger amount of solar energy than the second color and/or second material. Thus, rotating the balloon such that the first material is facing the sun may act to heat the envelope 302 as well as the gas inside the envelope 302. In this way, the buoyancy force of the envelope 302 may increase. By rotating the balloon such that the second material is facing the sun, the temperature of gas inside the envelope 302 may decrease. Accordingly, the buoyancy force may decrease. In this manner, the buoyancy force of the balloon could be adjusted by changing the temperature/volume of gas inside the envelope 302 using solar energy. In such embodiments, it is possible that a bladder 310 may not be a necessary element of balloon 300. Thus, various contemplated embodiments, altitude control of balloon 300 could be achieved, at least in part, by adjusting the rotation of the balloon with respect to the sun.

Further, the balloon 300 may include a navigation system (not shown). The navigation system may implement station-keeping functions to maintain position within and/or move to a position in accordance with a desired topology. In particular, the navigation system may use altitudinal wind data to determine altitudinal adjustments that result in the wind carrying the balloon in a desired direction and/or to a desired location. The altitude-control system may then make adjustments to the density of the balloon chamber in order to effectuate the determined altitudinal adjustments and cause the balloon to move laterally to the desired direction and/or to the desired location.

Alternatively, the altitudinal adjustments may be computed by a ground-based control system and communicated to the high-altitude balloon. As another alternative, the altitudinal



adjustments may be computed by a ground-based or satellite-based control system and communicated to the high-altitude balloon. Furthermore, in some embodiments, specific balloons in a heterogeneous balloon network may be configured to compute altitudinal adjustments for other balloons and transmit the adjustment commands to those other balloons.

As shown, the balloon **300** also includes a cut-down system **308**. The cut-down system **308** may be activated to separate the payload **306** from the rest of balloon **300**. This functionality may be utilized anytime the payload needs to be accessed on the ground, such as when it is time to remove balloon **300** from a balloon network, when maintenance is due on systems within payload **306**, and/or when power supply **326** needs to be recharged or replaced.

In an exemplary embodiment, the cut-down system **308** may include a connector, such as a balloon cord, connecting the payload **306** to the envelope **302** and a means for severing the connector (e.g., a shearing mechanism or an explosive bolt). In an example embodiment, the balloon cord, which may be nylon, is wrapped with a nichrome wire. A current could be passed through the nichrome wire to heat it and melt the cord, cutting the payload **306** away from the envelope **302**. Other types of cut-down systems and/or variations on the illustrated cut-down system **308** are possible as well.

In an alternative arrangement, a balloon may not include a cut-down system. In such an arrangement, the navigation system may be operable to navigate the balloon to a landing location, in the event the balloon needs to be removed from the network and/or accessed on the ground. Further, it is possible that a balloon may be self-sustaining, such that it theoretically does not need to be accessed on the ground. In yet other embodiments, balloons may be serviced in-flight by specific service balloons or another type of service aerostat or service aircraft.

As described, movement and locations of balloons can vary since winds in the stratosphere may affect the locations of the balloons in a differential manner. A balloon in an example network may be configured to change its horizontal position by adjusting its vertical position (i.e., altitude). For example, by adjusting its altitude, a balloon may be able to find winds that will carry the balloon horizontally (e.g., latitudinally and/or longitudinally) to a desired horizontal location. Wind speed and/or direction may vary with altitude, and since current wind velocities as well as weather forecasts are available, in principle, a balloon may be directed to a location by identifying an altitude at which a wind direction takes a balloon to along a desired trajectory. However, a balloon without other forms of propulsion may be constrained to follow the wind and there may not be a single altitude with winds taking the balloon along the desired trajectory. In addition, to control a fleet of balloons, movement of the balloons should occur from one location above the surface of the Earth to another in a predictable manner.

In some examples, determining a fleetplan for a fleet of balloons may require staging balloons of the fleet into places ahead of time such that the balloons will be able to arrive at a landmark on time. For instance, a balloon may need to follow a trajectory that does not contribute usefully to goals during a first phase and a second phase such that the balloon can reach a desired landmark at the end of a third phase. This kind of staging may be useful when goal landmarks are spread non-uniformly throughout a fleet's coverage area. For example, there may be large regions (e.g., relative to a balloon's maximum range during a phase) requiring few or no balloons, but which may need to be traversed in order to satisfy a goal number of balloons for other regions. Such situations arise naturally, for example, when goal distributions are propor-

tional to population density: there are a limited number of people in the oceans, few people on many parts of land, and many people in major cities.

In one case, a system may be configured to receive (or determine) a starting location of each balloon of a fleet of balloons and receive a sequence of coverage requirements for the region for a planning time period. As an example, the sequence of coverage requirements for the time period may specify a goal number of balloons for each of a plurality of landmarks within the region at various times  $T$  during the time period. The system may also be configured to divide the time period into phases based on distinct values of  $T$  in the sequence of coverage requirements: a first phase may be defined from a starting time of the time period to an earliest  $T$ ; a second phase may be defined between the end of the first phase and a second earliest  $T$ , etc. Additionally, a set of landmarks may be associated with the start and end of each phase: the start of the first phase may use the initial location of each of the balloons as landmarks; the end of the first phase and the start of the second phase may share the set of landmarks associated with the earliest  $T$ , etc.

Based on the received information, the system may be further configured to determine trajectories that satisfy the sequence of coverage requirements for the region. To determine the trajectories, the system may initially establish possible routes for balloons within each phase. For each landmark at the start of a phase and for each landmark at the end of the phase, the system may determine how close to the end-of-phase landmark that a balloon starting from the start-of-phase landmark could get by traveling during the phase.

FIG. 4 illustrates example trajectories for a balloon that is traveling from a plurality of starting landmarks to a landmark E. Specifically, FIG. 4 illustrates trajectories from each of landmark A, B, C, D, and F to an ending landmark E. As shown in FIG. 4, a region **400** has been divided into a plurality of subregions, and the landmarks A-L have been established at the center of each subregion. Note that in some examples coverage requirements may specify a goal number of balloons to be located at one or more locations within the region **400**, and the goal number of balloons may be added to the nearest landmark to determine the goal number of balloons for each landmark. For example, if an individual coverage requirement indicates that five balloons are desired at location Loc1 and seven balloons are desired at location Loc2, the desired number of balloons for landmark A may be determined to be twelve.

Additionally, although the landmarks A-L have been distributed uniformly throughout the region **400**, the example is not meant to be limiting. In other instances, landmarks may be non-uniformly distributed within the region **400**. For example, if a region covers the entire Earth, one or more oceans or countries may not include any landmarks.

The example trajectories may be determined based on estimated or predicted wind conditions at one or more altitudes. In the example of FIG. 4, the wind conditions are assumed to generally include winds flowing from left-to-right with additional variation in the vertical direction from time  $T=0$  to  $T=1$ . Therefore, it is assumed that a balloon could not flow from landmarks G-L to landmark E from time  $T=0$  to  $T=1$ , and for convenience, only trajectories from landmarks A, B, C, D, and F are shown in FIG. 4.

In one example, a backward planner may be used to determine a measure of how close to a target location that a balloon could reach if the balloon was starting at another location within a region. For example, given a target time and target location as well as an initial time, the backward planner may be configured to determine, for a grid of starting locations



within a region, a measure of how close to the target location a balloon could get, if the balloon was flying from initial time to target time.

The backward planner may consider a problem of minimizing a cost function associated with a balloon path. As an example, the cost function may be a measure of how close a balloon gets to a target location by following a path. FIG. 5 illustrates an example cost value map where location A is a starting location of a balloon, location B is a desired final location, and example cost values are assigned to cells surrounding location B. For a time period from  $T=0$  to  $T=N$ , the time period can be discretized in order to consider a situation of the balloon at times  $T=0, 1, 2, \dots, N$ . The region can also be discretized by dividing the region into a finite number of cells. For instance, FIG. 5 illustrates a region that has been divided into a number of cells (e.g., a  $10 \times 7$  grid of cells). Although the region is represented as a rectangular grid, the example is not meant to be limiting.

At time  $T=N$  (e.g., a final time of a planning phase), there is no planning to do since the balloon will be at a final location at this time, and a cost value can be assigned to each cell based on a final location of the balloon. For time  $T=N$ , cost values may be systematically assigned such that a cell in which the desired location resides has a low cost and cells further away may be assigned a higher cost value (e.g., the farther away from the desired location, the higher the cost value). Thus, as shown, a cell in which the desired location B resides can be assigned a cost value of zero, and cost values for surrounding cells can be assigned higher cost values based on a distance away from the desired location cell. Cost values for all cells can be populated in this manner by increasing the cost value for a given cell based on the distance from the desired location.

Following, cost value maps of all cells may be generated for each time step  $T=0, 1, 2, \dots, N-1$ . To construct the cost value maps, a map may be generated for time  $T=N-1$ , and maps for previous times may be generated working backwards, such as next generating a map for time  $T=N-2$ , followed by  $T=N-3$ , and so on.

After generating the map at time  $T=N$ , to generate a map for a next previous time step, possible altitudes and corresponding possible locations that can be reached over the time interval from  $T=N-1$  to  $T=N$  for a balloon present in every cell at time  $T=N-1$  can be estimated. For each cell and for one or more altitudes, an estimation of where the balloon will arrive by starting at that cell and flying at that altitude during the time interval can be determined. For example, weather predictions at time  $T=N-1$  may indicate that, for a given cell, different wind conditions exist for different altitudes. FIG. 6 illustrates an area divided into cells (e.g., a  $10$  by  $7$  grid of cells), and for a given cell labeled X, example wind vectors  $A_1$ - $A_5$  that are determined based on wind conditions at five different altitude levels.

In some examples, wind vectors associated with different altitudes may be the same for each cell. In other examples, wind vectors associated with different altitudes may vary based on the location of the cell. For each cell, an estimation of where a balloon would travel to (e.g., destination cell) at the end of the time period if the balloon started at the cell and moved with a given wind vector is determined. To do so, a direction and magnitude (e.g., speed) of the estimated wind vector can be used to estimate a trajectory and distance that the balloon travels, and to determine estimated destination cells (e.g.,  $A_{1(EST)}$ - $A_{5(EST)}$ ). FIG. 6 illustrates estimated destination cells for a balloon starting at cell X, and moving with one of the five wind vectors  $A_1$ - $A_5$ . Example simulations can

be performed to simulate movement of the balloons due to the estimated wind directions and speed, for example.

In some examples, further estimations may be made by a balloon traveling at a first altitude for a portion of a time interval and then changing to a second altitude for a remainder of the time interval. It could be possible that by traveling at any given number of combinations of altitudes, the winds may be able to carry the balloon to a desired location over the time interval.

As shown in the example in FIG. 6, wind vector  $A_3$  results in the balloon reaching a closest proximity to the desired final location B at time  $T=N$ . Cost values for the balloon traveling along the wind vectors  $A_1$ - $A_5$  may be determined based on the stored cost values at  $T=N$  for the destination cells  $A_{1(EST)}$ - $A_{5(EST)}$ . For example, the wind vector  $A_3$  may be assigned a value of zero, and cost values for a remainder of the wind vectors may be assigned higher cost values for causing the balloon to travel to a location further away from the desired destination (e.g., such as a  $A_4$  having cost of one,  $A_5$  having cost of four,  $A_2$  having cost of two, and  $A_1$  having cost of three).

An altitude that results in a smallest cost value can be determined, and the altitude and cost value can be stored for the cell X for time  $T=N-1$ . Thus, in the example described, for cell X at time  $T=N-1$ , a balloon in cell X that has a desired destination location B at time  $T=N$  can be instructed to travel at altitude  $A_3$ . The same determinations may be performed for each cell of the area.

The process above may be repeated to determine cost value maps for time interval  $T=N-2$  based on predicted wind conditions for  $T=N-2$  to  $T=N-1$  and the stored cost values for  $T=N-1$ . Further, the process may be repeated for time intervals  $T=N-3$ ,  $T=N-4$ , and so on until  $T=0$ . In some examples, each cell of a given map has estimates of directions that are based on estimates of directions as determined for a subsequent time interval. This enables planning for a balloon to take a given path at a first time interval so that the balloon will be at a location needed during the next time interval to take appropriate action, and so on.

Referring back to the example illustrated by FIG. 4, the backward planner may be able to determine, for each starting landmark A-L, a minimum distance from landmark E that a balloon could get, if the balloon was flying from  $T=0$  to  $T=1$ . For example, a cost value for each of landmarks A-L may be proportional to a distance such that a given cost value can be converted to a given three-dimensional distance. The backward planner may also be able to determine, for each starting landmark A-L, which altitude a balloon should fly at  $T=0$  in order to get as close as possible to the landmark E by  $T=1$ .

The backward planner is provided as one example mechanism for establishing possible routes between landmarks during a phase. However, other mechanisms that can determine, for a particular starting location/time and an ending location/time, how close to the ending location a balloon can get in the time period allotted, based on predicted wind velocities at one or more altitudes may also be utilized.

In some examples, filtering heuristics may also be applied to further refine the possible routes between landmarks during a phase. FIGS. 7A-7B illustrate examples of determining a set of starting landmarks from which a balloon could reach a landmark E.

In one example, thresholding may be used to impose a notion of how close to a target landmark that a balloon needs to be able to get in order to be considered to be capable of reaching the landmark. For example, a set of starting landmarks may be selected based on a comparison between the minimum distances associated with paths from landmarks



A-D and F and a threshold distance. Based on the trajectories shown in FIG. 7A, balloons traveling from landmarks A, B, and F could get within a threshold distance 702 of the landmark E. Therefore, the landmarks A, B, and F may be retained as a set of starting landmarks from which a balloon could reach the landmark E, while the landmarks D and C may be removed.

In another example, a set of starting landmarks may be selected based on a confidence of reaching a target landmark from a starting landmark and/or a cost value associated with flying to the target landmark from a starting landmark. As shown in FIG. 7B, each path from starting landmarks A-D and F includes an associated probability, as indicated by a percentage on a scale of 1 to 100. The probability may be determined based on the variability of one or more wind conditions (e.g., wind direction and/or speed). In one example, based on the associated probabilities, landmarks A and B and C may be selected as a set of possible starting landmarks because their associated probabilities are greater than a confidence threshold (e.g., 85%).

Additionally, as shown in FIG. 7B, each path from starting landmarks A-D and F includes an associated cost, as indicated by a dollar amount. In one example, the cost may be associated with a cost in power that is necessary to adjust the altitude of a balloon to reach one or more altitudes during the path. Accordingly, a set of starting landmarks may be selected based on a comparison between the cost values associated with each starting landmark and a cost threshold. For instance, starting landmark B may be selected as the only possible starting landmark because the path associated with landmark B has the only associated cost that is below 3\$.

In still another example, a set of starting landmarks may be determined from among the possible starting landmarks by retaining the top N number of starting landmarks (e.g., sorted to prefer minimum distances from within the target landmark that a balloon could get). It is contemplated that one or any combination of the described filtering heuristics may be used to determine a set of starting landmarks from which a target landmark may be reached.

FIG. 8 illustrates an example graph 800 of possible routes between a plurality of landmarks in accordance with an example embodiment. Continuing with the example from FIG. 4, landmarks A-F are provided. Each edge between two landmarks represents a trajectory a balloon could follow to get from a landmark at the beginning of a phase to a landmark at the end of the phase. For instance, a balloon at landmark A at time T=0 may be able to get to landmark D or E by time T=1. Additionally, a balloon at landmark A at time T=1 may be able to get to landmark D or E by time T=2. Note that although the edges for phase 1 are the same as the edges for phase 2 in FIG. 8, in other examples, edges may vary among phases.

Once the graph is constructed, an optimization method may be used to choose how many balloons flow across each connection in the graph (i.e., from each starting landmark to each ending landmark) during each phase. The optimization method may be configured to obey one or more consistency constraints. For example, a number of vehicles leaving a landmark in phase P may be set to be the same as the number of balloons that arrived at that landmark at the end of phase P-1. As another example, a number of vehicles at any landmark or traveling along any edge cannot be negative.

If the sequence of coverage requirements is treated as strict requirements, the problem of determining trajectories for the fleet may be considered a constraint satisfaction problem, where optimization means solving the constraint graph subject to the constraint that each goal is satisfied. If it is allow-

able to sometimes leave some goals unsatisfied, the problem may be considered an optimization problem: find the configuration that, for example, minimizes the number of unsatisfied goals, or some other measure of the goals' dissatisfaction.

Such a problem can be solved using a variety of methods, including as a linear program, as a minimum-cost-flow problem, via gradient descent, via stochastic optimization, or using other existing or future developed methods.

For instance, the problem may be solved as a linear program. In one case, the goal for the linear program is to take a time-varying description of how many balloons are desired at each landmark ("the goal distribution"), the initial location of each balloon, and a summary of which goal landmarks are reachable from which other landmarks at a given time, along with a measure of how accurately a balloon can be expected to reach a goal landmark, and determine a fleet-wide policy that maximizes a likelihood of satisfying the goal distribution. Note that the goal distribution is not concerned with determining where any particular balloon goes, but rather that the correct numbers of balloons are in the correct areas at each phase.

To solve the problem as a linear program, constants, variables, and constraints may be defined.

The constants may include:

startingAtOrigin[origin]=the initial distribution; the number of balloons at each location at the beginning of a planning period

requiredAtGoal[phase, goal]=the goal distribution; how many balloons desired at each goal at the end of each phase

score[phase, origin, goal]=accuracy estimate for the path from origin to goal during phase.

The variables may include:

A[phase, origin, goal]=how many vehicles move from origin to goal during phase; defined only if it is possible to reach the goal from the origin given the wind at that phase.

atGoal[phase, goal]=the number of balloons at goal at the end of phase

deficit[phase, goal]=the number of balloons by which the goal distribution was under-satisfied at goal at the of phase, that is:  $\max(0, \text{requiredAtGoal} - \text{atGoal})$ .

The constraints may include:

Non-negative flow

$\forall \text{origin}, \forall \text{goal}, \forall \text{phase}: 0 \leq A[\text{phase}, \text{origin}, \text{goal}]$

Conservation on the origin side

$\forall \text{origin}: \sum_{\text{goal}} A[0, \text{origin}, \text{goal}] = \text{startingAtOrigin}[\text{origin}]$

$\forall \text{origin}, \forall \text{phase} > 0: \sum_{\text{goal}} A[\text{phase}, \text{origin}, \text{goal}] = \text{atGoal}[\text{phase}-1, \text{origin}]$

Conservation on the goal side

$\forall \text{goal}, \forall \text{phase}: \text{atGoal}[\text{phase}, \text{goal}] = \sum_{\text{origin}} A[\text{phase}, \text{origin}, \text{goal}]$

Non-negative deficits

$\forall \text{goal}, \forall \text{phase}: 0 \leq \text{deficit}[\text{phase}, \text{goal}]$

$\forall \text{goal}, \forall \text{phase}: \text{requiredAtGoal}[\text{phase}, \text{goal}] \leq \text{atGoal}[\text{phase}, \text{goal}] + \text{deficit}[\text{phase}, \text{goal}]$

Based on the constants, variables, and constraints defined above, a linear program solver may be used to minimize the example objective function shown below:

$$\alpha \cdot c_{\text{error}} \cdot (\sum_{\text{phase}} \sum_{\text{goal}} \text{deficit}[\text{phase}, \text{goal}]) + (1 - \alpha) \cdot (\sum_{\text{phase}} \sum_{\text{origin}} \sum_{\text{goal}} \text{score}[\text{phase}, \text{origin}, \text{goal}] A[\text{phase}, \text{origin}, \text{goal}])$$

The objective function penalizes fleet plans that put less than the minimum number of balloons at some landmark; for each balloon fewer than desired, the penalty goes up. The



optimal plan, therefore, is the plan that minimizes the objective function. An additional term is also included that minimizes a measure of how close to a target landmark that each balloon can get. For instance, the additional term seeks to minimize the sum of the score constant for each path traversed. The additional term causes the linear program solver to prefer accurate trajectories and penalize trajectories that do not reliably deliver the balloon directly to a goal landmark.

In the example objective function,  $c_{error}$  is a weight parameter and  $\alpha$  may be another parameter that is used to adjust the relative importance of minimizing the number of missed goals versus preferring accurate trajectories. For instance, increasing  $\alpha \cdot c_{error}$  may cause the linear program solver to prefer fleet trajectories that satisfy a goal distribution better, at the cost of being willing to use less accurate flight paths. Also, increasing  $(1-\alpha)$  may cause the linear program solver to prefer fleet plans that use accurate/reliable flight paths, at the cost of being willing to leave more of a goal distribution unsatisfied.

In further examples, the score constant may be a measure of a confidence associated with whether a balloon traveling from origin could reach goal during phase. For instance, the measure of confidence may be a probability that is determined based on the variability of one or more wind conditions (e.g., wind direction and/or speed). In addition, the configuration of the linear program could further be modified such that the number of balloons atGoal is an expected number of balloons that is determined based on the probabilities of each balloon reaching the goal landmark. For example, if a balloon has a 75% chance of reaching a goal landmark, the balloons contribution to atGoal may be 0.75 balloons instead of 1 balloon.

In another variant, the score constant may be replaced or augmented with a measure of cost associated with a path from origin to goal during phase. For instance, the cost may be a cost-in-dollars, cost-in-power, or other cost measure. The linear program solver may then be configured to minimize total cost in the determined trajectories.

After the number of balloons to travel to each landmark during each phase has been solved for, a respective series of landmarks for each balloon to travel to over the period of time may be determined. In one example, for a first phase, each balloon may travel from its initial location to a first landmark as determined by the solution to the optimization problem. After considering which landmark that each balloon travels to during the first phase, the numbers of balloons that travel to each respective landmark may be used to determine which landmark that each balloon should travel to during the second phase.

An example determination of which landmark that each balloon should travel to is described with respect to FIGS. 9A and 9B. As shown in FIG. 9A, a solution to the optimization problem may indicate that, during a first phase, four balloons travel from a first landmark 902 to a second landmark 904 and five balloons travel from a third landmark 906 to the second landmark 904. The solution may further indicate that six balloons travel from the second landmark 904 to a fourth landmark 908 during the second phase and three balloons travel from the second landmark 904 to a fifth landmark 910 during the second phase. Thus, the number along each edge in FIG. 9A represents a number of balloons traveling along that edge.

In one example, based on the numbers of balloons traveling along each edge, six of the balloons that traveled to the second landmark 904 during the first phase may be assigned to travel to the fourth landmark 908 during the second phase and three of the balloons that traveled to the second landmark 904 during the first phase may be assigned to travel to the fifth

landmark 910 during the second phase. As shown in FIG. 9B, each arrow represents an individual balloon's trajectory. The example trajectories are not meant to be limiting and other configurations are also possible.

The determination may be repeatedly performed for each landmark and each subsequent phase, based on which balloons were assigned to travel to the landmark in a previous phase. As a result, a respective series of landmarks for each balloon to travel to may be found. In some examples, detailed trajectories for each balloon may also be determined by simulating the path of each balloon as the balloon travels to each landmark and follows one or more winds at one or more altitudes.

FIG. 10 illustrates example trajectories for each of a plurality of balloons. The plurality of balloons may travel through a region 1000 that is characterized by a plurality of landmarks. For each trajectory, the location of the corresponding balloon at  $T=1$ ,  $T=2$ , and  $T=3$  is noted. As shown in FIG. 10, a first balloon, that is initially located at Loc1, may be assigned to travel to landmarks D, G, and J. Based on the estimated winds for the region 1000, the first balloon may be expected to follow the path represented by the solid line. Similarly, a trajectory for a second balloon that is assigned to travel from Loc2 to landmark E, from landmark E to landmark H, and from landmark H to landmark K is shown. Additionally, a trajectory for a third balloon that is assigned to travel from Loc3 to landmark F, from landmark F to landmark I, and from landmark I to landmark K is shown.

It may also be desirable to evaluate or plan for the distribution of the balloons as each of the balloons pass between landmarks. In one case, an administrator might desire the fleet of balloons to not have any point (e.g., a point on the Earth, or a point within the stratosphere) to be more than a set distance from at least N balloons. This may allow the fleet of balloons to provide a desired sensor coverage, or to provide continuous wireless communication coverage for a region (or multiple regions). In another example, an administrator might wish to have every balloon be located within a set distance of at least N other balloons. Such a configuration may enable the fleet of balloons to maintain a balloon-to-balloon communication network. These types of immediate goals may be implemented using reactive, real-time control systems, such as flocking algorithms.

Flocking behavior may be described as the collective motion of a large number of self-propelled elements that is exhibited by many living beings, such as birds of a flock, fish, bacteria, and insects. Simulations, models, and algorithms that have been developed to emulate the flocking behaviors of living beings have also been applied to control the "flocking" behavior of other elements.

In some instances, the flocking behavior arises from rules that are followed by individual elements, in the absence of any central coordination. As an example, a flocking model may be governed by three rules: separation, alignment, and cohesion. Separation may cause each element to avoid crowding neighboring elements. Alignment may cause each element to steer towards an average heading of one or more neighboring elements. Cohesion may cause each element to steer towards an average position of one or more neighboring elements. Thus, based on the spatial relationship between an element and one or more neighboring elements and/or the behavior of one or more neighboring elements, each element may make a flocking-based control choice. FIG. 11 provides an example method that allows information from a long-term plan to be incorporated into a flocking-based control choice for each vehicle of a plurality of vehicles.



FIG. 11 is an example block diagram of a method 1100 to determine a flocking-based direction of travel. Method 1100 shown in FIG. 11 presents an embodiment of a method that, for example, could be used with the balloon networks 100 and 204, for example, and may be performed by a device, a server, a control station, a balloon, or a combination of any components of the balloon networks 100 and 204. Method 1100 may include one or more operations, functions, or actions as illustrated by one or more of blocks 1102-1108. Although the blocks are illustrated in a sequential order, these blocks may in some instances be performed in parallel, and/or in a different order than those described herein. Also, the various blocks may be combined into fewer blocks, divided into additional blocks, and/or removed based upon the desired implementation.

In addition, for the method 1100 and other processes and methods disclosed herein, the flowchart shows functionality and operation of one possible implementation of present embodiments. In this regard, each block may represent a module, a segment, or a portion of program code, which includes one or more instructions executable by a processor for implementing specific logical functions or steps in the process. The program code may be stored on any type of computer readable medium, for example, such as a storage device including a disk or hard drive. The computer readable medium may include a non-transitory computer readable medium, for example, such as computer-readable media that stores data for short periods of time like register memory, processor cache and Random Access Memory (RAM). The computer readable medium may also include non-transitory media, such as secondary or persistent long term storage, like read only memory (ROM), optical or magnetic disks, compact-disc read only memory (CD-ROM), for example. The computer readable media may also be any other volatile or non-volatile storage systems. The computer readable medium may be considered a computer readable storage medium, a tangible storage device, or other article of manufacture, for example.

In addition, for the method 1100 and other processes and methods disclosed herein, each block in FIG. 11 may represent circuitry that is wired to perform the specific logical functions in the process.

At block 1102, the method 1100 includes receiving a sequence of coverage requirements for a region and an associated period of time. An individual coverage requirement may indicate a desired distribution (e.g., a geographic distribution) of vehicles among the region at a respective time period within the period of time. In one example, each coverage requirement may be specified in the form of a goal distribution. For instance, the goal distribution may be determined based on population densities for different parts of the region or short-term coverage demands for one or more locations within the region.

At block 1104, the method 1100 includes, based on the sequence of coverage requirements, determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time. In one example, a sequence of intended destinations for each balloon of a fleet of balloons may be determined by receiving the sequence of intended destinations from a server or database. In another example, a linear objective function may be minimized to determine a number of vehicles to travel from each respective landmark to each additional landmark. The linear objective function may take into account possible routes between each landmark for each phase as well as other consistency constraints to determine the numbers of vehicles.

Based on the determined numbers of vehicles, a sequence of landmarks for each of one or more vehicles may be determined.

At block 1106, the method 1100 includes, based on a determined sequence of intended destinations for a vehicle of the one or more vehicles and based on a desired spatial relationship between the vehicle and one or more neighboring vehicles for a given time period, determining a flocking-based direction of travel for the vehicle for the given time period. In one example, a flocking algorithm may be modified to determine the flocking-based direction of travel for each vehicle based on a derived target distribution. In another example, the flocking algorithm may be modified to determine the flocking-based direction of travel for each vehicle based on a preferred direction of travel that is indicated by a control policy for each vehicle. In either example, the modified flocking algorithm determines flocking-based directions of travel that achieve immediate goals (e.g., a desired spatial relationship between a vehicle and one or more neighboring vehicles) while incorporating the staging and long-term planning associated with the sequence of coverage requirements.

As described above, a flocking algorithm may be governed by one or more flocking rules. Each of the rules may be considered to determine one or more components contributing to the determined flocking behavior (e.g., speed and direction) for a vehicle. FIGS. 12A and 12B are conceptual illustrations of flocking behaviors. As shown in FIGS. 12A and 12B, a flocking-based direction of travel for an individual vehicle 1200 may be determined based on a spatial relationship between the individual vehicle 1200 and one or more neighboring vehicles 1202A-D that are within a threshold distance of the individual vehicle 1200. The threshold distance may be a configurable parameter that controls which other vehicles the flocking behavior of the individual vehicle 1200 is dependent upon. In FIGS. 12A and 12B, the circle 1204 represents the boundary of the threshold distance. Thus, the flocking behavior of the individual vehicle 1200 may be dependent upon the behavior (e.g., position and/or velocity) of neighboring vehicles 1202A-D.

In some instances, it may be desirable to keep a specified amount of space between the individual vehicle 1200 and the neighboring vehicles 1202A-D. In order to do so, a determination may be made whether any of the neighboring vehicles 1202A-D are within a separation distance, represented by the dotted circle 1206, of the individual vehicle 1200. If any of the neighboring vehicles 1202A-D are within the separation distance, the individual vehicle 1200 may be caused to move in an opposite direction of the direction to the neighboring vehicle(s). As shown in FIG. 12A, as a result of the spatial relationship between the individual vehicle 1200 and the neighboring vehicle 1202C, the individual vehicle may be caused to move in a direction 1208 that is opposite to the direction towards the neighboring vehicle 1202C. The direction 1208 may be one component of the flocking-based direction of travel for the individual vehicle 1200.

In other instances, it may be desirable to keep the individual vehicle 1200 close to other vehicles. In order to do so, the position of each of the neighboring vehicles 1202A-D may be determined and averaged. The individual vehicle 1200 may then be caused to move towards the average position of the neighboring vehicles 1202A-D. As shown in FIG. 12B, the individual vehicle 1200 may be caused to move towards the average position 1210, and a direction 1212 towards the average position 1210 may be another component of the flocking-based direction of travel for the individual vehicle 1200.



Additionally, the individual vehicle **1200** may be caused to head in the same direction as the neighboring vehicles **1202A-D**. To do so, the average of each of the headings of the neighboring vehicles **1202A-D** may be determined. In FIG. **12B**, the average heading of the neighboring vehicles **1202A-D** is represented by the direction **1214**.

Any of the directions **1208**, **1212**, and **1214** may be combined to determine the flocking-based direction of travel for the individual vehicle **1200**. Further, each of the directions **1208**, **1212**, and **1214** may include an associated weighting parameter that impacts the relative importance of each of the directions **1208**, **1212**, and **1214** when the directions are combined to determine a net flocking-based direction.

The flocking behaviors described with reference to FIGS. **12A** and **12B** may be determined in real-time, as the vehicles are traveling throughout a region. For instance, each of the directions **1208**, **1212**, and **1214** may be re-determined periodically, at a fixed interval. However, it may also be desirable to incorporate long-term staging and planning goals into a flocking-based control choice that is made for each vehicle.

In one example, a flocking algorithm may be configured to incorporate a target distribution. For instance, a target distribution for various points in time may be determined based on the long-term plan for each vehicle. An example method for determining a target distribution is described with respect to FIGS. **13A** and **13B**.

A system may be configured to determine a detailed trajectory for each vehicle based on a respective series of landmarks that each vehicle is assigned to travel to over a period of time. Given the detailed trajectory for each vehicle, the system may then be able to determine a target distribution for any time period of the period of time. Initially, the location of each vehicle at the time period may be determined from the detailed trajectory for each vehicle. FIG. **13A** illustrates estimated locations of each vehicle of a fleet of vehicles. In FIG. **13A**, each X represents a location of a vehicle within a region **1300** at time  $T=2$ .

FIG. **13B** illustrates a derived target distribution **1302** that is determined based on the estimated locations of FIG. **13A**. In one example, to determine the target distribution **1302** for the time  $T=2$ , the region **1300** may be divided into a plurality of cells. Initially, the value of each cell may be set to zero (i.e., an empty distribution). For each vehicle, the distribution in a neighborhood around each vehicle's location may be incremented. For example, for a first vehicle **1304**, the value for the cell that the vehicle is located in, as well as cells above, below, to the left, and to the right of the cell may be increased by 0.2. Additionally, the process may be repeated for each additional vehicle within the region **1300**.

The contribution by each vehicle within the neighborhood around the vehicle may be uniform or non-uniform. For instance, the shape of the contribution is a tunable parameter. The spatial size of the neighborhood may also be a tunable parameter. Additionally, the total contribution from each vehicle may be set to one vehicle, or to a different amount. Mathematically, the process of determining a target distribution from each vehicle's location may be referred to as convolving each vehicle's location with an interest kernel. For instance, the interest kernel may be an isotropic Gaussian.

Some flocking algorithms may be capable of directly incorporating a target distribution for each time period. At each time period, the flocking algorithm may bias the direction of each vehicle such that the vehicles are arranged in a distribution that is similar to a target distribution. In other cases, parameters of the flocking algorithm may be adjusted based on values of the target distribution. FIG. **14** is another

example block diagram of a method **1400** to determine a flocking-based direction of travel.

At block **1402**, a series of landmarks for each vehicle to travel to may be determined. At block **1404**, the trajectory for each vehicle may be simulated based on the series of landmarks that each vehicle is expected to travel to. At block **1406**, a target distribution may be determined for a given time period. For each vehicle, a flocking-based direction may be determined for the given time period. To do so, a vehicle may be selected and a location of the vehicle may be determined at block **1408**. At block **1410**, a desired inter-vehicle distance and/or a desired number of neighboring vehicles may be determined by referencing the target distribution. For example, the desired inter-vehicle distance can be adjusted to prefer a shorter distance if the location of the vehicle is a location where the target distribution indicates a greater vehicle density is needed. In one case, the inter-vehicle distance may be inversely proportional to the vehicle density. Similarly, the desired number of neighboring vehicles may be selected by determining the density of the target distribution at a location of the vehicle. If the location is a location where the target distribution indicates a greater vehicle density is needed, the number of neighboring vehicles may be increased.

At block **1412**, a flocking-based direction may be determined. In some examples, the desired inter-vehicle distance and/or number of neighboring vehicles may be used to adjust parameters of the flocking algorithm. For example, the inter-vehicle distance may be the separation distance described in FIG. **12A**. Further, to increase the number of neighboring vehicles, the threshold distance that governs which other vehicles that an individual vehicle's behavior is dependent upon may be increased and/or the separation distance may be decreased.

In other examples, a flocking algorithm may be configured to receive a preferred direction of travel and bias a vehicle to travel in the preferred direction. FIG. **15** is still another example block diagram of a method **1500** to determine a flocking-based direction of travel. Initially, at block **1502**, a series of landmarks for each vehicle to travel to may be determined. Subsequently, at block **1504**, a control policy and/or fitness values may be determined for each vehicle.

A system may derive a control policy for each vehicle to satisfy its series of landmarks to travel to. For example, the control policy may specify, for each location within a region, the preferred direction of travel in order to get as close as possible to the next intended landmark or as close as possible to as many of the remaining landmarks in a series of landmarks as possible. Optionally the control policy may vary over time. For instance, at a first time period, the control policy may indicate the preferred direction of travel in order to get as close as possible to a first landmark while at a second time period, the control policy may indicate a different preferred direction of travel in order to get as close as possible to a next landmark. The backward planner described with respect to FIGS. **5** and **6** may be used to determine a preferred direction of travel for each location of a grid of locations of a region based on a desired target location.

The system may also derive a fitness value for each location rather than or in addition to an explicit control policy. For example, the fitness value for a location may be proportional to a distance from a goal landmark that a vehicle could reach by starting from the location. The optimal policy for a vehicle at a particular location may then be the action among the currently available actions that leads to the location with the best fitness value.



At block **1506**, an individual vehicle may be selected and the location of the vehicle may be determined. Then, at block **1508**, a preferred direction of travel for the vehicle may be determined. For instance, the preferred direction of travel may be determined by looking up the preferred direction of travel in the control policy for the vehicle based on the location of the vehicle. Alternatively, the preferred direction of travel may be found by determining the direction of travel that leads the vehicle to a location having a best fitness value (e.g., a minimum fitness value).

In some examples, rather than using just a vehicle's location, the control policy for the vehicle may also be encoded in terms of a vehicle's extended configuration. For example, the extended configuration may include the vehicle's location as well as other status information (e.g., remaining fuel reserves, age, health, etc.). In such an example, the vehicle's configuration may then be used to look up the preferred direction of travel in the control policy.

At block **1510**, a flocking-based direction of travel for the vehicle may be determined. For instance, the flocking algorithm may receive as input the preferred direction of travel. The preferred direction of travel may bias the flocking algorithm to determine a flocking-based direction that moves the vehicle in the preferred direction of travel. For instance, the preferred direction of travel may be a component direction that is combined with other component directions (e.g., components that are determined based on the spatial relationship between the individual vehicle and one or more neighboring vehicles and/or based on the behavior of one or more neighboring vehicles) to determine a net flocking-based direction.

In an example in which fitness values are determined for each vehicle and for each location of a region, the fitness values may also allow the flocking algorithm to consider how sub-optimal a different direction would be. For instance, a fitness value associated with a direction in which the flocking algorithm proposes that the vehicle heads may indicate that heading in the proposed direction may cause the vehicle to miss an intended goal landmark. In some examples, the flocking algorithm may be configured to only allow vehicles to deviate from an optimal direction of travel when the deviation's fitness value meets certain criteria (e.g., a difference between an optimal fitness value and a fitness value associated with the deviation is less than a threshold).

Any of the methods **1100**, **1400**, and **1500** may further include providing an instruction to each vehicle to travel in the flocking-based direction of travel. In an example in which vehicles of the plurality of vehicles include a balloon within a data network, each vehicle may be instructed to fly at a specific altitude. In other examples, an instruction can be sent to the balloon to indicate to change an altitude, and the instruction may be sent when the altitude change is needed.

In some examples, any of the methods **1100**, **1400**, and **1500** or portions of the methods **1100**, **1400**, and **1500** may be performed by a ground-based station (e.g., the ground-based station **112** in FIG. **1**) or a central system (e.g., the central control system **200** in FIG. **2**), and instructions can be provided to individual vehicles. In other examples, the methods **1100**, **1400**, and **1500** or portions of the methods **1100**, **1400**, and **1500** may be performed by the vehicles themselves, or by processors or computing devices residing on the vehicle, for example. The vehicles may receive any necessary information for performing the methods **1100**, **1400**, or **1500** from a server or ground-base station or from other vehicles. In further examples, any of the methods **1100**, **1400**, and **1500** may be performed by a combination of the vehicles and by ground-based stations, and each may be in communication so as to perform functions of the method **1100**.

In some instances, the methods **1100**, **1400**, and **1500** or a portion of the methods **1100**, **1400**, and **1500** may be subsequently re-executed to determine subsequent flocking-based directions of travel. For instance, the methods **1100**, **1400**, and **1500** may be performed at a new time period for any vehicles of a fleet of vehicles as the spatial relationship between each vehicle and one or more neighboring vehicles changes.

It should be understood that arrangements described herein are for purposes of example only. As such, those skilled in the art will appreciate that other arrangements and other elements (e.g. machines, interfaces, functions, orders, and groupings of functions, etc.) can be used instead, and some elements may be omitted altogether according to the desired results. Further, many of the elements that are described are functional entities that may be implemented as discrete or distributed components or in conjunction with other components, in any suitable combination and location.

While various aspects and embodiments have been disclosed herein, other aspects and embodiments will be apparent to those skilled in the art. The various aspects and embodiments disclosed herein are for purposes of illustration and are not intended to be limiting, with the true scope being indicated by the following claims, along with the full scope of equivalents to which such claims are entitled. It is also to be understood that the terminology used herein is for the purpose of describing particular embodiments only, and is not intended to be limiting.

What is claimed is:

**1.** A method comprising:

receiving a sequence of coverage requirements for a region and an associated period of time, wherein the region is characterized by a plurality of landmarks, and wherein individual coverage requirements of the sequence of coverage requirements indicate a desired number of vehicles for different landmarks of the plurality of landmarks at a respective time period within the period of time;

determining, by a processor, a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time to satisfy the sequence of coverage requirements; and determining, by the processor, a flocking-based direction of travel for a vehicle of the one or more vehicles and a given time period based on: (a) a determined sequence of intended destinations for the vehicle and (b) a desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

**2.** The method of claim **1**, further comprising instructing the vehicle to travel in the flocking-based direction.

**3.** The method of claim **1**, further comprising: determining a target geographic distribution of vehicles for the given time period using a determined sequence of intended destinations for each of a plurality of vehicles of the fleet of vehicles; and

determining the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period based on a location of the vehicle and a density of the target geographic distribution for the location.

**4.** The method of claim **3**, wherein determining the target geographic distribution for the given time period using a determined sequence of intended destinations for each of a plurality of vehicles comprises:

determining a trajectory for each vehicle over the period of time using the determined sequence of intended destinations for the vehicle;



25

determining an estimated location of each vehicle for the given time period based on the determined trajectory for each vehicle; and

determining the target geographic distribution based on the estimated location of each vehicle for the given time period.

5. The method of claim 1, wherein the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period comprises a desired minimum inter-vehicle distance between the vehicle and each of the one or more neighboring vehicles.

6. The method of claim 1, wherein the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period comprises a desired minimum number of vehicles to be located within a threshold distance of the vehicle.

7. The method of claim 1, further comprising:

based on the determined sequence of intended destinations for the vehicle, determining a control policy for the vehicle, wherein the control policy for the vehicle indicates a direction of travel for the vehicle based on a location of the vehicle; and

determining the flocking-based direction of travel for the vehicle for the given time period based on the control policy for the vehicle and the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

8. The method of claim 7, wherein vehicles of the fleet of vehicles include a balloon within a data network that is operable to provide data communication via optical or radio-frequency links, and the method further comprises:

determining the control policy for the vehicle based on estimated winds at one or more altitudes that the vehicle could follow.

9. The method of claim 1, wherein the region is divided into a plurality of cells, and the method further comprises:

based on the determined sequence of intended destinations for the vehicle, determining respective fitness values for each cell of the plurality of cells, wherein a respective fitness value is based at least in part on a distance from an intended destination that the vehicle could reach by starting from a given cell of the plurality of cells;

determining the flocking-based direction of travel for the vehicle for the given time period based on the determined fitness values and the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

10. The method of claim 9, further comprising determining the flocking-based direction of travel for the vehicle for the given time period based at least in part on a direction to a cell having a minimum fitness value and on the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

11. The method of claim 1, wherein vehicles of the fleet of vehicles include a balloon within a data network that is operable to provide data communication via optical or radio-frequency links.

12. The method of claim 1, wherein determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time to satisfy the sequence of coverage requirements comprises:

determining for each of one or more phases during the period of time, a number of vehicles to travel between one or more respective landmarks of the plurality of landmarks and one or more additional landmarks of the

26

plurality of landmarks during the phase to satisfy the sequence of coverage requirements; and  
determining the respective sequence of intended destinations for each of the one or more vehicles of the fleet of vehicles to travel to over the period of time based on the determined number of vehicles for the one or more phases.

13. A non-transitory computer-readable storage medium having stored therein instructions, that when executed by a computing device, cause the computing device to perform functions comprising:

receiving a sequence of coverage requirements for a region and an associated period of time, wherein the region is characterized by a plurality of landmarks, and wherein individual coverage requirements of the sequence of coverage requirements indicate a desired number of vehicles for different landmarks of the plurality of landmarks at a respective time period within the period of time;

determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time to satisfy the sequence of coverage requirements; and

determining a flocking-based direction of travel for a vehicle of the one or more vehicles and a given time period based on: (a) a determined sequence of intended destinations for the vehicle and (b) a desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

14. The non-transitory computer-readable storage medium of claim 13, wherein the functions further comprise:

determining a target geographic distribution of vehicles for the given time period using the determined sequence of intended destinations for each of a plurality of vehicles of the fleet of vehicles; and

determining the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period based on a location of the vehicle and a density of the target geographic distribution for the location.

15. The non-transitory computer-readable storage medium of claim 13, wherein the functions further comprise:

based on the determined sequence of intended destinations for the vehicle, determining a control policy for the vehicle, wherein the control policy for the vehicle indicates a direction of travel for the vehicle based on a location of the vehicle; and

determining the flocking-based direction of travel for the vehicle for the given time period based on the control policy for the vehicle and the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

16. The non-transitory computer-readable storage medium of claim 13, wherein the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period comprises one or more of: a desired minimum inter-balloon distance between the vehicle and each of the one or more neighboring vehicles; and a desired minimum number of vehicles to be located within a threshold distance of the vehicle.

17. A system, comprising:

at least one processor; and

data storage comprising program instructions executable by the at least one processor to cause the system to perform functions comprising:

receiving a sequence of coverage requirements for a region and an associated period of time, wherein the

27

region is characterized by a plurality of landmarks, and wherein individual coverage requirements of the sequence of coverage requirements indicate a desired number of vehicles for different landmarks of the plurality of landmarks at a respective time period 5 within the period of time;

determining a respective sequence of intended destinations for each of one or more vehicles of a fleet of vehicles to travel to over the period of time to satisfy the sequence of coverage requirements; and 10

determining a flocking-based direction of travel for a vehicle of the one or more vehicles for the and a given time period based on: (a) a determined sequence of intended destinations vehicle and (b) a desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period. 15

**18.** The system of claim **17**, wherein the functions further comprise:

determining a target geographic distribution of vehicles for the given time period using the determined sequence of intended destinations for each of a plurality of vehicles of the fleet of vehicles; and 20

28

determining the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period based on a location of the vehicle and a density of the target geographic distribution for the location.

**19.** The system of claim **17**, wherein the functions further comprise:

based on the determined sequence of intended destinations for the vehicle, determining a control policy for the vehicle, wherein the control policy for the vehicle indicates a direction of travel for the vehicle based on a location of the vehicle; and

determining the flocking-based direction of travel for the vehicle for the given time period based on the control policy for the vehicle and the desired spatial relationship between the vehicle and one or more neighboring vehicles for the given time period.

**20.** The system of claim **17**, wherein vehicles of the fleet of vehicles include a balloon within a data network that is operable to provide data communication via optical or radio-frequency links.

\* \* \* \* \*



UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 8,781,727 B1  
APPLICATION NO. : 13/742209  
DATED : January 15, 2013  
INVENTOR(S) : Keith Bonawitz and Dan Piponi

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Claims

Col. 27, claim 17, line 13, delete the words “for the”;

Col. 27, claim 17, line 15, insert the words -- for the -- between “destinations” and “vehicles”.

Signed and Sealed this  
Seventh Day of October, 2014



Michelle K. Lee  
*Deputy Director of the United States Patent and Trademark Office*